

Veritas Storage Foundation™ Cluster File System High Availability Administrator's Guide

Solaris

6.0

Veritas Storage Foundation™ Cluster File System High Availability Administrator's Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 6.0

Document version: 6.0.4

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4
Section 1 Introducing Storage Foundation Cluster File System High Availability	31
Chapter 1 Overview of Storage Foundation Cluster File System High Availability	32
About Veritas Storage Foundation Cluster File System High Availability	32
About Veritas File System	33
Logging	33
Extents	34
File system disk layouts	34
About Veritas Volume Manager	34
About Veritas Dynamic Multi-Pathing	35
About Veritas Operations Manager	35
About Storage Foundation Cluster File System High Availability solutions	36
Chapter 2 How Storage Foundation Cluster File System High Availability works	38
How Storage Foundation Cluster File System High Availability works	39
When to use Storage Foundation Cluster File System High Availability	40
Storage Foundation Cluster File System High Availability architecture	41
About the symmetric architecture	41
About Storage Foundation Cluster File System High Availability primary/secondary failover	42
About single-host file system semantics using Group Lock Manager	42

About Veritas File System features supported in cluster file systems	42
Veritas File System features in cluster file systems	43
Veritas File System features not in cluster file systems	45
About Storage Foundation Cluster File System High Availability and the Group Lock Manager	45
Storage Foundation Cluster File System High Availability namespace	46
About asymmetric mounts	46
Primary and secondary	47
Determining or moving primaryship	47
Synchronize time on Cluster File Systems	48
File system tunables	48
Setting the number of parallel fsck threads	48
Storage Checkpoints	49
Storage Foundation Cluster File System High Availability backup strategies	49
Parallel I/O	50
I/O error handling policy	51
Recovering for I/O failures	51
Single network link and reliability	52
Configuring a low-priority link	52
Split-brain and jeopardy handling	53
About I/O fencing	54
About I/O fencing for SFCFSHA in virtual machines that do not support SCSI-3 PR	55
About preventing data corruption with I/O fencing	55
About I/O fencing components	56
About I/O fencing configuration files	60
How I/O fencing works in different event scenarios	62
About server-based I/O fencing	66
About secure communication between the SFCFSHA cluster and CP server	70
Storage Foundation Cluster File System High Availability and Veritas Volume Manager cluster functionality agents	74
Veritas Volume Manager cluster functionality	74
 Chapter 3 How Veritas File System works	 75
Veritas File System features	75
Extent-based allocation	77
Extent attributes	78
Fast file system recovery	78

Extended mount options	79
Enhanced data integrity modes	80
Enhanced performance mode	81
Temporary file system mode	81
Improved synchronous writes	81
Support for large files and large file systems	81
Access Control Lists	82
Storage Checkpoints	82
FileSnaps	82
Online backup	82
Quotas	83
Support for databases	83
Cluster file systems	84
Cross-platform data sharing	84
File Change Log	85
Reverse path name lookup	85
Multi-volume file system support	85
SmartTier	85
Thin Reclamation of a file system	86
Partitioned directories	86
Data deduplication on Solaris SPARC	86
File compression	87
Veritas File System performance enhancements	87
Enhanced I/O performance	88
Delayed allocation for extending writes	89
Using Veritas File System	89
Online system administration	89
Application program interface	91

Chapter 4	How Veritas Volume Manager works	92
	How VxVM works with the operating system	92
	How data is stored	93
	How VxVM handles storage management	93
	Physical objects	94
	Virtual objects	96
	About the configuration daemon in VxVM	103
	Multiple paths to disk arrays	104
	Volume layouts in VxVM	104
	Non-layered volumes	105
	Layered volumes	105
	Layout methods	105
	Concatenation, spanning, and carving	106

Striping (RAID-0)	108
Mirroring (RAID-1)	111
Striping plus mirroring (mirrored-stripe or RAID-0+1)	112
Mirroring plus striping (striped-mirror, RAID-1+0 or RAID-10)	113
RAID-5 (striping with parity)	114
Online relay layout	121
How online relay layout works	121
Limitations of online relay layout	124
Transformation characteristics	125
Transformations and volume length	125
Volume resynchronization	126
Dirty flags	126
Resynchronization process	126
Hot-relocation	127
Dirty region logging	127
Log subdisks and plexes	128
Sequential DRL	128
SmartSync recovery accelerator	128
Volume snapshots	130
Comparison of snapshot features	131
FastResync	133
How FastResync works	133
How non-persistent FastResync works with snapshots	134
How persistent FastResync works with snapshots	135
DCO volume versioning	138
Effect of growing a volume on the FastResync map	140
FastResync limitations	141
Volume sets	142
Configuration of volumes on SAN storage	142

Chapter 5	How CVM works	145
	About the cluster functionality of VxVM	145
	Overview of clustering	146
	Overview of cluster volume management	147
	Private and shared disk groups	149
	Activation modes of shared disk groups	150
	CVM resiliency to storage connectivity failures	152
	Effect of disk connectivity on cluster reconfiguration	157
	Limitations of shared disk groups	157
	CVM initialization and configuration	158
	Cluster reconfiguration	158
	Volume reconfiguration	161

	Node shutdown	164
	Cluster shutdown	165
	Dirty region logging in cluster environments	165
	How DRL works in a cluster environment	166
	Multiple host failover configurations	167
	Import lock	167
	Failover	167
	Corruption of disk group configuration	168
Chapter 6	How Veritas Dynamic Multi-Pathing works	170
	How DMP works	170
	Device discovery	174
	How DMP monitors I/O on paths	176
	Load balancing	178
	Disabling Sun Multipath IO (MPxIO)	179
	Dynamic Reconfiguration	179
	About booting from DMP devices	180
	DMP in a clustered environment	180
	Veritas Volume Manager co-existence with Oracle Automatic Storage Management (ASM) disks	181
Section 2	Provisioning storage	184
Chapter 7	Provisioning new storage	185
	Provisioning new storage	185
	Growing the existing storage by adding a new LUN	186
	Growing the existing storage by growing the LUN	186
	Displaying SFCFSA information with vxlist	187
Chapter 8	Advanced allocation methods for configuring storage	188
	Customizing allocation behavior	188
	Setting default values for vxassist	189
	Using rules and persistent attributes to make volume allocation more efficient	192
	Creating volumes of a specific layout	198
	Types of volume layouts	198
	Creating a mirrored volume	199
	Creating a striped volume	201
	Creating a RAID-5 volume	203
	Creating a volume on specific disks	204

Chapter 9

Creating volumes on specific media types	206
Specifying ordered allocation of storage to volumes	206
Site-based allocation	209
Changing the read policy for mirrored volumes	209
Chapter 9 Creating and mounting VxFS file systems	211
Creating a VxFS file system	211
Block size	212
Intent log size	213
Example of creating a file system	213
Converting a file system to VxFS	214
Example of converting a file system	215
Mounting a VxFS file system	215
log mount option	216
delaylog mount option	216
tmplog mount option	216
logiosize mount option	217
nodatainlog mount option	218
blkclear mount option	218
mincache mount option	218
convosync mount option	219
ioerror mount option	220
largefiles nolargefiles mount option	222
cio mount option	223
mntlock mount option	224
ckptautomnt mount option	224
Combining mount command options	224
Example of mounting a file system	225
Unmounting a file system	225
Example of unmounting a file system	225
Resizing a file system	226
Extending a file system using fsadm	226
Shrinking a file system	227
Reorganizing a file system	228
Displaying information on mounted file systems	229
Example of displaying information on mounted file systems	230
Identifying file system types	230
Example of determining a file system's type	231
Monitoring free space	231
Monitoring fragmentation	231
Thin Reclamation	233

Chapter 10	Extent attributes	235
	About extent attributes	235
	Reservation: preallocating space to a file	236
	Fixed extent size	236
	How the fixed extent size works with the shared extents	237
	Other extent attribute controls	237
	Commands related to extent attributes	239
	Example of setting an extent attribute	239
	Example of getting an extent attribute	240
	Failure to preserve extent attributes	240
Section 3	Administering multi-pathing with DMP	242
Chapter 11	Administering Dynamic Multi-Pathing	243
	Discovering and configuring newly added disk devices	243
	Partial device discovery	244
	Discovering disks and dynamically adding disk arrays	245
	Third-party driver coexistence	247
	How to administer the Device Discovery Layer	249
	Making devices invisible to VxVM	262
	Making devices visible to VxVM	264
	About enabling and disabling I/O for controllers and storage processors	264
	About displaying Veritas Dynamic Multi-Pathing (DMP) information	265
	Displaying the paths to a disk	266
	Administering DMP using vxddmadm	269
	Retrieving information about a DMP node	271
	Displaying consolidated information about the DMP nodes	272
	Displaying the members of a LUN group	273
	Displaying paths controlled by a DMP node, controller, enclosure, or array port	274
	Displaying information about controllers	276
	Displaying information about enclosures	277
	Displaying information about array ports	278
	Displaying DMP path information for devices under third-party driver control	279
	Displaying extended device attributes	280
	Suppressing or including devices from VxVM control	282
	Gathering and displaying DMP I/O statistics	283

Setting the attributes of the paths to an enclosure	289
Displaying the path redundancy level of a device or enclosure under DMP control	290
Specifying the minimum number of active paths for a device or enclosure under DMP control	290
Displaying the DMP I/O policy	291
Specifying the DMP I/O policy	292
Disabling I/O for paths, controllers or array ports	298
Enabling I/O for paths, controllers or array ports	299
Renaming an enclosure	300
Configuring DMP's response to I/O failures	301
Configuring the DMP I/O throttling mechanism	302
Configuring DMP's Low Impact Path Probing	303
Configuring DMP's Subpaths Failover Groups (SFG)	304
Displaying recovery option values	304
Configuring DMP path restoration policies	305
Stopping the DMP path restoration thread	307
Displaying the status of the DMP path restoration thread	307
Configuring array policy modules	307
DMP coexistence with native multi-pathing	308

Chapter 12 Dynamic reconfiguration of devices 310

About online dynamic reconfiguration	310
Reconfiguring a LUN online that is under DMP control	311
Removing LUNs dynamically from an existing target ID	311
Adding new LUNs dynamically to a new target ID	314
About detecting target ID reuse if the operating system device tree is not cleaned up	315
Scanning an operating system device tree after adding or removing LUNs	316
Cleaning up the operating system device tree after removing LUNs	316
Dynamic LUN expansion	318
Replacing a host bus adapter on an M5000 server managed by DMP	319
Upgrading the array controller firmware online	323

Chapter 13 Managing devices 325

Displaying disk information	325
Displaying disk information with vxdiskadm	326
Changing the disk device naming scheme	327
Displaying the disk-naming scheme	328

	Setting customized names for DMP nodes	329
	Regenerating persistent device names	330
	Changing device naming for TPD-controlled enclosures	330
	Simple or nopriv disks with enclosure-based naming	332
	About the Array Volume Identifier (AVID) attribute	333
	About disk installation and formatting	335
	Adding and removing disks	336
	Adding a disk to VxVM	336
	Removing disks	345
	Renaming a disk	348
Chapter 14	Event monitoring	349
	About the Dynamic Multi-Pathing (DMP) event source daemon (vxesd)	349
	Fabric Monitoring and proactive error detection	350
	Dynamic Multi-Pathing (DMP) automated device discovery on Solaris	351
	Dynamic Multi-Pathing (DMP) discovery of iSCSI and SAN Fibre Channel topology	352
	DMP event logging	352
	Starting and stopping the Dynamic Multi-Pathing (DMP) event source daemon	353
Section 4	Administering Storage Foundation Cluster File System High Availability and its components	354
Chapter 15	Administering Storage Foundation Cluster File System High Availability and its components	355
	About Storage Foundation Cluster File System High Availability administration	355
	Administering CFS	356
	Adding CFS file systems to VCS configuration	356
	Using cfsmount to mount CFS file systems	357
	Resizing CFS file systems	357
	Verifying the status of CFS file systems	357
	Verifying CFS port	358
	CFS agents and AMF support	358
	CFS agent log files	358

Storage Foundation Cluster File System High Availability	
commands	358
mount, fsclustadm, and fsadm commands	359
Time synchronization for Cluster File Systems	360
Growing a Storage Foundation Cluster File System High	
Availability	361
The /etc/vfstab file	361
When the Storage Foundation Cluster File System High Availability	
primary fails	361
Storage Checkpoints on Storage Foundation Cluster File System	
High Availability	362
Snapshots on Storage Foundation Cluster File System High	
Availability	362
Administering VCS	364
Viewing available Veritas device drivers	365
Configuring VCS to start Oracle with a specified Pfile	365
Verifying VCS configuration	365
Starting and stopping VCS	365
Configuring destination-based load balancing for LLT	366
Administering CVM	366
Listing all the CVM shared disks	366
Establishing CVM cluster membership manually	366
Methods to control CVM master selection	367
About setting cluster node preferences for master failover	367
About changing the CVM master manually	372
Importing a shared disk group manually	376
Deporting a shared disk group manually	376
Starting shared volumes manually	376
Evaluating the state of CVM ports	376
Verifying if CVM is running in an SFCFSHA cluster	377
Verifying CVM membership state	377
Verifying the state of CVM shared disk groups	378
Verifying the activation mode	378
CVM log files	378
Requesting node status and discovering the master node	379
Determining if a LUN is in a shareable disk group	380
Listing shared disk groups	380
Creating a shared disk group	381
Importing disk groups as shared	382
Converting a disk group from shared to private	383
Moving objects between shared disk groups	383
Splitting shared disk groups	383
Joining shared disk groups	384

Changing the activation mode on a shared disk group	384
Setting the detach policy for shared disk groups	384
Handling cloned disks in a shared disk group	385
Creating volumes with exclusive open access by a node	385
Setting exclusive open access to a volume by a node	385
Displaying the cluster protocol version	386
Displaying the supported cluster protocol version range	386
Recovering volumes in shared disk groups	387
Obtaining cluster performance statistics	387
Administering CVM from the slave node	388
Administering ODM	390
Verifying the ODM port	390
Starting ODM	390
Administering I/O Fencing	390
About administering I/O fencing	390
About the vxfsentsthdw utility	391
About the vxfsenadm utility	399
About the vxfsenclearpre utility	404
About the vxfsenswap utility	407
About administering the coordination point server	418
About migrating between disk-based and server-based fencing configurations	437
Enabling or disabling the preferred fencing policy	453
About VXFEN tunable parameters	455
Administering SFCFSHA global clusters	459
About setting up a disaster recovery fire drill	459
About configuring the fire drill service group using the Fire Drill Setup wizard	460
Verifying a successful fire drill	462
Scheduling a fire drill	463

Chapter 16	Administering sites and remote mirrors	464
	About sites and remote mirrors	464
	About site-based allocation	467
	About site consistency	468
	About site tags	469
	About the site read policy	469
	Making an existing disk group site consistent	470
	Configuring a new disk group as a Remote Mirror configuration	471
	Fire drill — testing the configuration	472
	Simulating site failure	473
	Verifying the secondary site	473

	Recovery from simulated site failure	473
	Changing the site name	473
	Resetting the site name for a host	474
	Administering the Remote Mirror configuration	474
	Configuring site tagging for disks or enclosures	474
	Configuring automatic site tagging for a disk group	475
	Configuring site consistency on a volume	476
	Examples of storage allocation by specifying sites	476
	Displaying site information	478
	Failure and recovery scenarios	479
	Recovering from a loss of site connectivity	480
	Recovering from host failure	481
	Recovering from storage failure	481
	Recovering from site failure	481
	Recovering from disruption of connectivity to storage at the remote sites from hosts on all sites	482
	Recovering from disruption to connectivity to storage at all sites from the hosts at a site	482
	Automatic site reattachment	482
Section 5	Optimizing I/O performance	484
Chapter 17	Veritas File System I/O	485
	About Veritas File System I/O	485
	Buffered and Direct I/O	486
	Direct I/O	486
	Unbuffered I/O	487
	Data synchronous I/O	487
	Concurrent I/O	488
	Cache advisories	489
	Freezing and thawing a file system	489
	Getting the I/O size	490
	About Storage Foundation and High Availability Solutions products database accelerators	490
Chapter 18	Veritas Volume Manager I/O	491
	Veritas Volume Manager throttling of administrative I/O	491

Section 6	Veritas Extension for Oracle Disk Manager	493
Chapter 19	Using Veritas Extension for Oracle Disk Manager	494
	About Oracle Disk Manager	494
	How Oracle Disk Manager improves database performance	496
	About Oracle Disk Manager and Veritas Storage Foundation Cluster File System High Availability	497
	About Oracle Disk Manager and Oracle Managed Files	498
	How Oracle Disk Manager works with Oracle Managed Files	498
	Setting up Veritas Extension for Oracle Disk Manager	500
	Preparing existing database storage for Oracle Disk Manager	501
	Converting Quick I/O files to Oracle Disk Manager files	501
	Verifying that Oracle Disk Manager is configured	502
	Disabling the Oracle Disk Manager feature	504
	Using Cached ODM	505
	Enabling Cached ODM for file systems	506
	Modifying Cached ODM settings for individual files	506
	Adding Cached ODM settings via the cachemap	507
	Making the caching settings persistent across mounts	508
Section 7	Clustered NFS	509
Chapter 20	Clustered NFS	510
	About Clustered NFS	510
	Requirements for Clustered NFS	510
	Understanding how Clustered NFS works	511
	Basic design	511
	Internal Clustered NFS functionality	511
	cfsshare manual page	513
	Configure and unconfigure Clustered NFS	513
	Configure Clustered NFS	513
	Unconfiguring Clustered NFS	515
	Reconciling major and minor numbers for NFS shared disks	515
	Checking the major and minor number for VxVM volumes	515
	Administering Clustered NFS	518
	Displaying the NFS shared CFS file systems	518
	Sharing a CFS file system previously added to VCS	518
	Unsharing the previous shared CFS file system	519

	Adding an NFS shared CFS file system to VCS	519
	Deleting the NFS shared CFS file system from VCS	520
	Adding a virtual IP address to VCS	520
	Deleting a virtual IP address from VCS	520
	Adding an IPv6 virtual IP address to VCS	521
	Deleting an IPv6 virtual IP address from VCS	521
	Changing the share options associated with an NFS share	521
	Sharing a file system checkpoint	521
	Samples for configuring a Clustered NFS	522
	Sample main.cf file	526
	How to mount an NFS-exported file system on the NFS clients	532
	Disabling SMF for NFS daemons on Solaris 10	533
	Debugging Clustered NFS	535
Section 8	Common Internet File System	536
Chapter 21	Common Internet File System	537
	About Common Internet File System	537
	Requirements for Common Internet File System	538
	Understanding how Samba works	538
	Configuring Clustered NFS and Common Internet File System on	
	CFS	538
	cfsshare manual page	538
	Configuring Common Internet File System in user mode	539
	Configuring Common Internet File System in domain mode	541
	Configuring Common Internet File System in ads mode	542
	Administering Common Internet File System	544
	Sharing a CFS file system previously added to VCS	546
	Unsharing the previous shared CFS file system	547
	Sample main.cf file	547
	Debugging Common Internet File System	553
Section 9	Using Point-in-time copies	554
Chapter 22	Understanding point-in-time copy methods	555
	About point-in-time copies	555
	When to use point-in-time copies	556
	Implementing point-in time copy solutions on a primary host	557
	Implementing off-host point-in-time copy solutions	559
	About Storage Foundation point-in-time copy technologies	565
	Comparison of Point-in-time copy solutions	566

Volume-level snapshots	567
Persistent FastResync of volume snapshots	567
Data integrity in volume snapshots	568
Third-mirror break-off snapshots	568
Space-optimized instant volume snapshots	569
Choices for snapshot resynchronization	570
Disk group split/join	570
Storage Checkpoints	571
How Storage Checkpoints differ from snapshots	572
How a Storage Checkpoint works	572
About Storage Rollbacks	576
Storage Checkpoints and Storage Rollback process	576
Types of Storage Checkpoints	578
About FileSnaps	580
Properties of FileSnaps	580
Concurrent I/O to FileSnaps	581
Copy-on-write and FileSnaps	581
Reading from FileSnaps	582
Block map fragmentation and FileSnaps	582
Backup and FileSnaps	582
About snapshot file systems	583
How a snapshot file system works	583

Chapter 23	Administering volume snapshots	585
	About volume snapshots	585
	How traditional third-mirror break-off snapshots work	586
	How full-sized instant snapshots work	587
	Linked break-off snapshot volumes	589
	Cascaded snapshots	590
	Creating a snapshot of a snapshot	591
	Creating multiple snapshots	593
	Restoring the original volume from a snapshot	594
	Creating instant snapshots	595
	Adding an instant snap DCO and DCO volume	596
	Creating and managing space-optimized instant snapshots	601
	Creating and managing full-sized instant snapshots	604
	Creating and managing third-mirror break-off snapshots	606
	Creating and managing linked break-off snapshot volumes	609
	Creating multiple instant snapshots	611
	Creating instant snapshots of volume sets	612
	Adding snapshot mirrors to a volume	614
	Removing a snapshot mirror	615

Removing a linked break-off snapshot volume	615
Adding a snapshot to a cascaded snapshot hierarchy	615
Refreshing an instant space-optimized snapshot	616
Reattaching an instant full-sized or plex break-off snapshot	616
Reattaching a linked break-off snapshot volume	617
Restoring a volume from an instant space-optimized snapshot	618
Dissociating an instant snapshot	618
Removing an instant snapshot	619
Splitting an instant snapshot hierarchy	619
Displaying instant snapshot information	620
Controlling instant snapshot synchronization	622
Listing the snapshots created on a cache	624
Tuning the autogrow attributes of a cache	624
Monitoring and displaying cache usage	625
Growing and shrinking a cache	625
Removing a cache	626
Creating traditional third-mirror break-off snapshots	626
Converting a plex into a snapshot plex	630
Creating multiple snapshots with the vxassist command	631
Reattaching a snapshot volume	632
Adding plexes to a snapshot volume	633
Dissociating a snapshot volume	633
Displaying snapshot information	634
Adding a version 0 DCO and DCO volume	635
Specifying storage for version 0 DCO plexes	636
Removing a version 0 DCO and DCO volume	638
Reattaching a version 0 DCO and DCO volume	638

Chapter 24	Administering Storage Checkpoints	640
	About Storage Checkpoints	640
	Storage Checkpoint administration	641
	Creating a Storage Checkpoint	642
	Removing a Storage Checkpoint	643
	Accessing a Storage Checkpoint	643
	Converting a data Storage Checkpoint to a nodata Storage Checkpoint	645
	Enabling and disabling Storage Checkpoint visibility	653
	Storage Checkpoint space management considerations	654
	Restoring from a Storage Checkpoint	654
	Examples of restoring a file from a Storage Checkpoint	655
	Storage Checkpoint quotas	660

Chapter 25	Administering FileSnaps	661
	FileSnap creation	661
	FileSnap creation over Network File System	661
	Using FileSnaps	662
	Using FileSnaps to create point-in-time copies of files	663
	Using FileSnaps to provision virtual desktops	663
	Using FileSnaps to optimize write intensive applications for virtual machines	664
	Using FileSnaps to create multiple copies of data instantly	664
	Comparison of the logical size output of the fsadm -S shared, du, and df commands	665
Chapter 26	Administering snapshot file systems	667
	Snapshot file system backups	667
	Snapshot file system performance	668
	About snapshot file system disk structure	668
	Differences between snapshots and Storage Checkpoints	669
	Creating a snapshot file system	670
	Examples of creating snapshot file systems	671
Section 10	Optimizing thin storage with Storage Foundation Cluster File System High Availability	672
Chapter 27	Understanding thin storage solutions in Storage Foundation Cluster File System High Availability	673
	About thin provisioning	673
	About thin optimization solutions in Storage Foundation Cluster File System High Availability	674
	About SmartMove	675
	SmartMove for thin provisioning	676
	About the Thin Reclamation feature	676
	Determining when to reclaim space on a thin reclamation LUN	677
	How automatic reclamation works	678

Chapter 28	Migrating data from thick storage to thin storage	679
	About using SmartMove to migrate to Thin Storage	679
	Migrating to thin provisioning	679
Chapter 29	Maintaining Thin Storage with Thin Reclamation	683
	Reclamation of storage on thin reclamation arrays	683
	Thin Reclamation of a disk, a disk group, or an enclosure	684
	Thin Reclamation of a file system	685
	Identifying thin and thin reclamation LUNs	685
	Displaying VxFS file system usage on thin reclamation LUNs	686
	Reclaiming space on a file system	689
	Reclaiming space on a disk, disk group, or enclosure	689
	Monitoring Thin Reclamation using the vxtask command	691
	Configuring automatic reclamation	692
Section 11	Maximizing storage utilization	694
Chapter 30	Understanding storage tiering with SmartTier	695
	About SmartTier	695
	SmartTier building blocks	696
	About VxFS multi-volume file systems	697
	About VxVM volume sets	697
	About volume tags	698
	How SmartTier works	698
	SmartTier file management	699
	SmartTier sub-file object management	699
	How the SmartTier policy works with the shared extents	700
	SmartTier in a High Availability (HA) environment	700
Chapter 31	Creating and administering volume sets	701
	About volume sets	701
	Creating a volume set	702
	Adding a volume to a volume set	703
	Removing a volume from a volume set	703
	Listing details of volume sets	703
	Stopping and starting volume sets	704
	Managing raw device nodes of component volumes	705
	Enabling raw device access when creating a volume set	706

	Displaying the raw device access settings for a volume set	707
	Controlling raw device access for an existing volume set	707
Chapter 32	Multi-volume file systems	709
	About multi-volume file systems	709
	About volume types	710
	Features implemented using multi-volume support	710
	Volume availability	711
	Creating multi-volume file systems	712
	Example of creating a multi-volume file system	712
	Converting a single volume file system to a multi-volume file system	713
	Adding a volume to and removing a volume from a multi-volume file system	715
	Adding a volume to a multi-volume file system	715
	Removing a volume from a multi-volume file system	715
	Forcibly removing a volume in a multi-volume file system	715
	Moving volume 0 in a multi-volume file system	716
	Volume encapsulation	716
	Encapsulating a volume	716
	Deencapsulating a volume	718
	Reporting file extents	718
	Examples of reporting file extents	719
	Load balancing	720
	Defining and assigning a load balancing allocation policy	721
	Rebalancing extents	721
	Converting a multi-volume file system to a single volume file system	722
Chapter 33	Administering SmartTier	724
	About SmartTier	724
	Supported SmartTier document type definitions	726
	Placement classes	726
	Tagging volumes as placement classes	727
	Listing placement classes	727
	Administering placement policies	728
	Assigning a placement policy	728
	Unassigning a placement policy	728
	Analyzing the space impact of enforcing a placement policy	729
	Querying which files will be affected by enforcing a placement policy	729
	Enforcing a placement policy	729

Validating a placement policy	731
File placement policy grammar	731
File placement policy rules	732
SELECT statement	732
CREATE statement	735
RELOCATE statement	737
DELETE statement	751
Calculating I/O temperature and access temperature	753
Multiple criteria in file placement policy rule statements	757
Multiple file selection criteria in SELECT statement clauses	757
Multiple placement classes in <ON> clauses of CREATE statements and in <TO> clauses of RELOCATE statements	758
Multiple placement classes in <FROM> clauses of RELOCATE and DELETE statements	759
Multiple conditions in <WHEN> clauses of RELOCATE and DELETE statements	759
File placement policy rule and statement ordering	760
File placement policies and extending files	762
Using SmartTier with solid state disks	762
Fine grain temperatures with solid state disks	763
Prefer mechanism with solid state disks	764
Average I/O activity with solid state disks	764
Frequent SmartTier scans with solid state disks	765
Quick identification of cold files with solid state disks	765
Example placement policy when using solid state disks	766

Chapter 34	Administering hot-relocation	771
	About hot-relocation	771
	How hot-relocation works	772
	Partial disk failure mail messages	775
	Complete disk failure mail messages	776
	How space is chosen for relocation	777
	Configuring a system for hot-relocation	778
	Displaying spare disk information	779
	Marking a disk as a hot-relocation spare	779
	Removing a disk from use as a hot-relocation spare	780
	Excluding a disk from hot-relocation use	781
	Making a disk available for hot-relocation use	782
	Configuring hot-relocation to use only spare disks	782
	Moving relocated subdisks	783
	Moving relocated subdisks using vxunreloc	783

	Restarting vxunreloc after errors	786
	Modifying the behavior of hot-relocation	786
Chapter 35	Deduplicating data on Solaris SPARC	788
	About deduplicating data	788
	About deduplication chunk size	789
	Deduplication and file system performance	790
	About the deduplication scheduler	790
	Deduplicating data	791
	Enabling and disabling deduplication on a file system	792
	Scheduling deduplication of a file system	792
	Performing a deduplication dry run	793
	Querying the deduplication status of a file system	794
	Starting and stopping the deduplication scheduler daemon	794
	Example of deduplicating a file system	795
	Deduplication results	796
	Deduplication supportability	797
	Deduplication use cases	797
	Deduplication limitations	797
Chapter 36	Compressing files	799
	About compressing files	799
	About the compressed file format	800
	About the file compression attributes	800
	About the file compression block size	801
	Compressing files with the vxcompress command	801
	Examples of using the vxcompress command	802
	Interaction of compressed files and other commands	803
	Interaction of compressed files and other features	804
	Interaction of compressed files and applications	805
	Use cases for compressing files	806
	Compressed files and databases	806
	Compressing all files that meet the specified criteria	810
Section 12	Reference	811
Chapter 37	Managing storage	812
	Moving volumes or disks	812
	Moving volumes from a VM disk	812
	Moving disks between disk groups	813
	Reorganizing the contents of disk groups	814

Monitoring and controlling tasks	827
Specifying task tags	827
Managing tasks with vxtask	828
Using vxnotify to monitor configuration changes	831
Performing online relayout	831
Permitted relayout transformations	832
Specifying a non-default layout	835
Specifying a plex for relayout	836
Tagging a relayout operation	836
Viewing the status of a relayout	836
Controlling the progress of a relayout	837
Adding a mirror to a volume	838
Mirroring all volumes	838
Mirroring volumes on a VM disk	839
Additional mirroring considerations	840
Configuring SmartMove	841
Removing a mirror	842
Setting tags on volumes	843
Managing disk groups	844
Disk group versions	844
Managing volumes	849
Managing plexes and subdisks	850
Reattaching plexes	850
Plex synchronization	853
Decommissioning storage	854
Removing a volume	854
Removing a disk from VxVM control	855
About shredding data	855
Shredding a VxVM disk	856
Failed disk shred operation results in a disk with no label	858
Removing and replacing disks	858

Chapter 38	Rootability	866
	Encapsulating a disk	866
	Failure of disk encapsulation	870
	Using nopriv disks for encapsulation	870
	Rootability	872
	Booting root volumes	873
	Boot-time volume restrictions	873
	Creating redundancy for the root disk	874
	Creating an archived back-up root disk for disaster recovery	875
	Mirroring an encapsulated root disk	875

	Booting from alternate boot disks	876
	Mirroring other file systems on the root disk	878
	Encapsulating SAN disks	879
	Administering an encapsulated boot disk	880
	Creating a snapshot of an encapsulated boot disk	881
	Growing an encapsulated boot disk	881
	Unencapsulating the root disk	883
Chapter 39	Managing tunable parameters	884
	Reference tables of Storage Foundation tunable parameters	884
	DMP tunable parameters	884
	Tunable parameters for VxVM	892
	How to change tunable parameters	906
	Methods to change Veritas Dynamic Multi-Pathing tunable parameters	906
	Methods to change Veritas Volume Manager tunable parameters	914
Chapter 40	Quotas	923
	About quota limits	923
	About quota files on Veritas File System	924
	About quota commands	924
	About quota checking with Veritas File System	926
	Using quotas	926
	Turning on quotas	926
	Turning on quotas at mount time	927
	Editing user and group quotas	927
	Modifying time limits	928
	Viewing disk quotas and usage	928
	Displaying blocks owned by users or groups	929
	Turning off quotas	929
Chapter 41	File Change Log	930
	About File Change Log	930
	About the File Change Log file	931
	File Change Log administrative interface	931
	File Change Log programmatic interface	934
	Summary of API functions	936
Chapter 42	Reverse path name lookup	937
	Reverse path name lookup	937

Appendix A	Disk layout	939
	About disk layouts	939
	VxFS Version 7 disk layout	941
	VxFS Version 8 disk layout	942
	VxFS Version 9 disk layout	942
	Using UNIX Commands on File Systems Larger than One TB	943
Appendix B	Command reference	944
	Veritas Volume Manager command reference	944
	CVM commands supported for executing on the slave node	964
	Veritas Volume Manager manual pages	971
	Section 1M — administrative commands	972
	Section 4 — file formats	975
	Section 7 — device driver interfaces	976
	Veritas File System command summary	976
	Veritas File System manual pages	979
Appendix C	Creating a starter database	986
	Creating a database for Oracle 11gr2	986
	Creating database tablespace on shared raw VxVM volumes	986
Index		989

Introducing Storage Foundation Cluster File System High Availability

- [Chapter 1. Overview of Storage Foundation Cluster File System High Availability](#)
- [Chapter 2. How Storage Foundation Cluster File System High Availability works](#)
- [Chapter 3. How Veritas File System works](#)
- [Chapter 4. How Veritas Volume Manager works](#)
- [Chapter 5. How CVM works](#)
- [Chapter 6. How Veritas Dynamic Multi-Pathing works](#)

Overview of Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [About Veritas Storage Foundation Cluster File System High Availability](#)
- [About Veritas File System](#)
- [About Veritas Volume Manager](#)
- [About Veritas Dynamic Multi-Pathing](#)
- [About Veritas Operations Manager](#)
- [About Storage Foundation Cluster File System High Availability solutions](#)

About Veritas Storage Foundation Cluster File System High Availability

Veritas Storage Foundation Cluster File System High Availability by Symantec extends Veritas Storage Foundation to support shared data in a storage area network (SAN) environment. Using Storage Foundation Cluster File System High Availability, multiple servers can concurrently access shared storage and files transparently to applications.

Veritas Storage Foundation Cluster File System High Availability also provides increased automation and intelligent management of availability and performance.

Veritas Storage Foundation Cluster File System High Availability includes Veritas Cluster Server, which adds high availability functionality to the product.

Before you install the product, read the *Veritas Storage Foundation Cluster File System High Availability Release Notes*.

To install the product, follow the instructions in the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

For information on high availability environments, read the Veritas Cluster Server documentation.

About Veritas File System

A file system is simply a method for storing and organizing computer files and the data they contain to make it easy to find and access them. More formally, a file system is a set of abstract data types (such as metadata) that are implemented for the storage, hierarchical organization, manipulation, navigation, access, and retrieval of data.

Veritas File System (VxFS) was the first commercial journaling file system. With journaling, metadata changes are first written to a log (or journal) then to disk. Since changes do not need to be written in multiple places, throughput is much faster as the metadata is written asynchronously.

VxFS is also an extent-based, intent logging file system. VxFS is designed for use in operating environments that require high performance and availability and deal with large amounts of data.

VxFS major components include:

- [Logging](#)
- [Extents](#)
- [File system disk layouts](#)

Logging

A key aspect of any file system is how to recover if a system crash occurs. Earlier methods required a time-consuming scan of the entire file system. A better solution is the method of logging (or journaling) the metadata of files.

VxFS logs new attribute information into a reserved area of the file system, whenever file system changes occur. The file system writes the actual data to disk only after the write of the metadata to the log is complete. If and when a system crash occurs, the system recovery code analyzes the metadata log and tries to clean up only

those files. Without logging, a file system check (`fsck`) must look at all of the metadata.

Intent logging minimizes system downtime after abnormal shutdowns by logging file system transactions. When the system is halted unexpectedly, this log can be replayed and outstanding transactions can be completed. The recovery time for file systems can be reduced to a few seconds, regardless of the file system size.

By default, VxFS file systems log file transactions before they are committed to disk, reducing time spent recovering file systems after the system is halted unexpectedly.

Extents

An extent is a contiguous area of storage in a computer file system, reserved for a file. When starting to write to a file, a whole extent is allocated. When writing to the file again, the data continues where the previous write left off. This reduces or eliminates file fragmentation.

Since VxFS is an extent-based file system, addressing is done through extents (which can consist of multiple blocks) rather than in single-block segments. Extents can therefore enhance file system throughput.

File system disk layouts

The disk layout is the way file system information is stored on disk. On VxFS, several disk layout versions, numbered 1 through 9, were created to support various new features and specific UNIX environments. Currently, only the Version 7, 8, and 9 disk layouts can be created and mounted. The Version 6 disk layout can be mounted, but only for upgrading to a supported version. No other versions can be created or mounted.

About Veritas Volume Manager

Veritas™ Volume Manager (VxVM) by Symantec is a storage management subsystem that allows you to manage physical disks and logical unit numbers (LUNs) as logical devices called volumes. A VxVM volume appears to applications and the operating system as a physical device on which file systems, databases and other managed data objects can be configured.

VxVM provides easy-to-use online disk storage management for computing environments and Storage Area Network (SAN) environments. By supporting the Redundant Array of Independent Disks (RAID) model, VxVM can be configured to protect against disk and hardware failure, and to increase I/O throughput.

Additionally, VxVM provides features that enhance fault tolerance and fast recovery from disk failure or storage array failure.

VxVM overcomes restrictions imposed by hardware disk devices and by LUNs by providing a logical volume management layer. This allows volumes to span multiple disks and LUNs.

VxVM provides the tools to improve performance and ensure data availability and integrity. You can also use VxVM to dynamically configure storage while the system is active.

About Veritas Dynamic Multi-Pathing

Veritas Dynamic Multi-Pathing (DMP) provides multi-pathing functionality for the operating system native devices configured on the system. DMP creates DMP metadevices (also known as DMP nodes) to represent all the device paths to the same physical LUN.

DMP is available as a component of Storage Foundation. DMP supports Veritas Volume Manager (VxVM) volumes on DMP metadevices, and Veritas File System (VxFS) file systems on those volumes.

DMP is also available as a stand-alone product, which extends DMP metadevices to support ZFS. You can create ZFS pools on DMP metadevices. DMP supports only non-root ZFS file systems.

Veritas Dynamic Multi-Pathing can be licensed separately from Storage Foundation products. Veritas Volume Manager and Veritas File System functionality is not provided with a DMP license.

DMP functionality is available with a Storage Foundation Enterprise license, SF HA Enterprise license, and Standard license.

Veritas Volume Manager (VxVM) volumes and disk groups can co-exist with ZFS pools, but each device can only support one of the types. If a disk has a VxVM label, then the disk is not available to ZFS. Similarly, if a disk is in use by ZFS, then the disk is not available to VxVM.

About Veritas Operations Manager

Symantec recommends use of Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

Veritas Operations Manager provides a centralized management console for Veritas Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports.

You can download Veritas Operations Manager at no charge at <http://go.symantec.com/vom>.

Refer to the Veritas Operations Manager documentation for installation, upgrade, and configuration instructions.

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. If you want to continue using VEA, a software version is available for download from

<http://www.symantec.com/operations-manager/support>. Veritas Storage Foundation Management Server is deprecated.

If you want to manage a single cluster using Cluster Manager (Java Console), a version is available for download from

https://www4.symantec.com/Vrt/offer?a_id=89446. You cannot manage the new features of this release using the Java Console. Veritas Cluster Server Management Console is deprecated.

About Storage Foundation Cluster File System High Availability solutions

Storage Foundation Cluster File System High Availability components and features can be used individually and in concert to improve performance, resilience, and ease of management for your storage and applications. Storage Foundation Cluster File System High Availability features can be used for:

- Improving database performance: you can use Storage Foundation Cluster File System High Availability database accelerators to improve I/O performance. SFHA Solutions database accelerators achieve the speed of raw disk while retaining the management features and convenience of a file system.
- Optimizing thin array usage: you can use Storage Foundation Cluster File System High Availability thin provisioning and thin reclamation solutions to set up and maintain thin storage.
- Backing up and recovering data: you can use Storage Foundation Cluster File System High Availability Flashsnap, Storage Checkpoints, and NetBackup point-in-time copy methods to back up and recover your data.
- Processing data off-host: you can avoid performance loss to your production hosts by using Storage Foundation Cluster File System High Availability volume snapshots.
- Optimizing test and development environments: you can optimize copies of your production database for test, decision modeling, and development purposes

using Storage Foundation Cluster File System High Availability point-in-time copy methods.

- Optimizing virtual desktop environments: you can use Storage Foundation Cluster File System High Availability FileSnap to optimize your virtual desktop environment.
- Maximizing storage utilization: you can use Storage Foundation Cluster File System High Availability SmartTier to move data to storage tiers based on age, priority, and access rate criteria.
- Migrating your data: you can use Storage Foundation Cluster File System High Availability Portable Data Containers to easily and reliably migrate data from one environment to another.

For a supplemental guide which documents Storage Foundation Cluster File System High Availability use case solutions using example scenarios:

See the *Veritas Storage Foundation™ and High Availability Solutions Guide*.

How Storage Foundation Cluster File System High Availability works

This chapter includes the following topics:

- [How Storage Foundation Cluster File System High Availability works](#)
- [When to use Storage Foundation Cluster File System High Availability](#)
- [Storage Foundation Cluster File System High Availability architecture](#)
- [About Veritas File System features supported in cluster file systems](#)
- [About Storage Foundation Cluster File System High Availability and the Group Lock Manager](#)
- [Storage Foundation Cluster File System High Availability namespace](#)
- [About asymmetric mounts](#)
- [Primary and secondary](#)
- [Determining or moving primaryship](#)
- [Synchronize time on Cluster File Systems](#)
- [File system tunables](#)
- [Setting the number of parallel fsck threads](#)
- [Storage Checkpoints](#)
- [Storage Foundation Cluster File System High Availability backup strategies](#)

- [Parallel I/O](#)
- [I/O error handling policy](#)
- [Recovering for I/O failures](#)
- [Single network link and reliability](#)
- [Split-brain and jeopardy handling](#)
- [About I/O fencing](#)
- [Storage Foundation Cluster File System High Availability and Veritas Volume Manager cluster functionality agents](#)
- [Veritas Volume Manager cluster functionality](#)

How Storage Foundation Cluster File System High Availability works

SFCFSHA simplifies or eliminates system administration tasks that result from the following hardware limitations:

- The SFCFSHA single file system image administrative model simplifies administration by enabling the execution of all file system management commands from any node.
- Because all servers in a cluster have access to SFCFSHA cluster-shareable file systems, keeping data consistent across multiple servers is automatic. All cluster nodes have access to the same data, and all data is accessible by all servers using single server file system semantics.
- Because all files can be accessed by all servers, applications can be allocated to servers to balance load or meet other operational requirements. Similarly, failover becomes more flexible because it is not constrained by data accessibility.
- Because each SFCFSHA file system can be on any node in the cluster, the file system recovery portion of failover time in an n -node cluster can be reduced by a factor of n by distributing the file systems uniformly across cluster nodes.
- Enterprise RAID subsystems can be used more effectively because all of their capacity can be mounted by all servers, and allocated by using administrative operations instead of hardware reconfigurations.
- Larger volumes with wider striping improve application I/O load balancing. Not only is the I/O load of each server spread across storage resources, but with SFCFSHA shared file systems, the loads of all servers are balanced against each other.

- Extending clusters by adding servers is easier because each new server's storage configuration does not need to be set up—new servers simply adopt the cluster-wide volume and file system configuration.
- The clusterized Oracle Disk Manager (ODM) feature that makes file-based databases perform as well as raw partition-based databases is available to applications running in a cluster.

When to use Storage Foundation Cluster File System High Availability

You should use SFCFSHA for any application that requires the sharing of files, such as for home directories and boot server files, Web pages, and for cluster-ready applications. SFCFSHA is also applicable when you want highly available standby data, in predominantly read-only environments where you just need to access data, or when you do not want to rely on NFS for file sharing.

Almost all applications can benefit from SFCFSHA. Applications that are not “cluster-aware” can operate on and access data from anywhere in a cluster. If multiple cluster applications running on different servers are accessing data in a cluster file system, overall system I/O performance improves due to the load balancing effect of having one cluster file system on a separate underlying volume. This is automatic; no tuning or other administrative action is required.

Many applications consist of multiple concurrent threads of execution that could run on different servers if they had a way to coordinate their data accesses. SFCFSHA provides this coordination. Such applications can be made cluster-aware allowing their instances to co-operate to balance client and data access load, and thereby scale beyond the capacity of any single server. In such applications, SFCFSHA provides shared data access, enabling application-level load balancing across cluster nodes.

SFCFSHA provides the following features:

- For single-host applications that must be continuously available, SFCFSHA can reduce application failover time because it provides an already-running file system environment in which an application can restart after a server failure.
- For parallel applications, such as distributed database management systems and Web servers, SFCFSHA provides shared data to all application instances concurrently. SFCFSHA also allows these applications to grow by the addition of servers, and improves their availability by enabling them to redistribute load in the event of server failure simply by reassigning network addresses.

- For workflow applications, such as video production, in which very large files are passed from station to station, the SFCFSHA eliminates time consuming and error prone data copying by making files available at all stations.
- For backup, the SFCFSHA can reduce the impact on operations by running on a separate server, accessing data in cluster-shareable file systems.

The following are examples of applications and how they might work with SFCFSHA:

- Using Storage Foundation Cluster File System High Availability on file servers
Two or more servers connected in a cluster configuration (that is, connected to the same clients and the same storage) serve separate file systems. If one of the servers fails, the other recognizes the failure, recovers, assumes the primaryship, and begins responding to clients using the failed server's IP addresses.
- Using Storage Foundation Cluster File System High Availability on web servers
Web servers are particularly suitable to shared clustering because their application is typically read-only. Moreover, with a client load balancing front end, a Web server cluster's capacity can be expanded by adding a server and another copy of the site. A SFCFSHA-based cluster greatly simplifies scaling and administration for this type of application.

Storage Foundation Cluster File System High Availability architecture

The Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) allows clustered servers to mount and use a file system simultaneously as if all applications using the file system were running on the same server. The Veritas Volume Manager cluster functionality (CVM) makes logical volumes and raw device applications accessible throughout a cluster.

This section includes the following topics:

- About the symmetric architecture
- About SFCFSHA primary/secondary failover
- About single-host file system semantics using Group Lock Manager

About the symmetric architecture

SFCFSHA uses a symmetric architecture in which all nodes in the cluster can simultaneously function as metadata servers. SFCFSHA still has some remnants of the old master/slave or primary/secondary concept. The first server to mount each cluster file system becomes its primary; all other nodes in the cluster become

secondaries. Applications access the user data in files directly from the server on which they are running. Each SFCFSHA node has its own intent log. File system operations, such as allocating or deleting files, can originate from any node in the cluster.

About Storage Foundation Cluster File System High Availability primary/secondary failover

If the server on which the SFCFSHA primary is running fails, the remaining cluster nodes elect a new primary. The new primary reads the intent log of the old primary and completes any metadata updates that were in process at the time of the failure.

If a server on which an SFCFSHA secondary is running fails, the primary reads the intent log of the failed secondary and completes any metadata updates that were in process at the time of the failure.

See [“When the Storage Foundation Cluster File System High Availability primary fails”](#) on page 361.

About single-host file system semantics using Group Lock Manager

SFCFSHA uses the Veritas Group Lock Manager (GLM) to reproduce UNIX single-host file system semantics in clusters. This is most important in write behavior. UNIX file systems make writes appear to be atomic. This means that when an application writes a stream of data to a file, any subsequent application that reads from the same area of the file retrieves the new data, even if it has been cached by the file system and not yet written to disk. Applications can never retrieve stale data, or partial results from a previous write.

To reproduce single-host write semantics, system caches must be kept coherent and each must instantly reflect any updates to cached data, regardless of the cluster node from which they originate. GLM locks a file so that no other node in the cluster can update it simultaneously, or read it before the update is complete.

About Veritas File System features supported in cluster file systems

The Veritas Storage Foundation Cluster File System High Availability is based on the Veritas File System (VxFS).

Most of the major features of VxFS local file systems are available on cluster file systems, including the following features:

- Extent-based space management that maps files up to a terabyte in size

- Fast recovery from system crashes using the intent log to track recent file system metadata updates
- Online administration that allows file systems to be extended and defragmented while they are in use

Every VxFS manual page has a section on Storage Foundation Cluster File System High Availability Issues with information on whether the command functions on a cluster-mounted file system and indicates any difference in behavior from local mounted file systems.

Veritas File System features in cluster file systems

[Table 2-1](#) describes the VxFS supported features and commands for SFCFSHA.

Table 2-1 Veritas File System features in cluster file systems

Features	Description
Compression	Compressing files reduces the space used by files, while retaining the accessibility of the files and being transparent to applications. See “About compressing files” on page 799.
Deduplication	You can perform post-process periodic deduplication in a file system to eliminate duplicate data without any continuous cost. See “About deduplicating data” on page 788.
Defragmentation	Free resources are initially aligned and allocated to files in an order that provides optimal performance. See “Defragmentation” on page 90.
Disk layout versions	SFCFSHA supports only disk layout Version 7, 8 and 9. Use the <code>fstyp -v special_device</code> command to ascertain the disk layout version of a VxFS file system. Use the <code>vxupgrade</code> command for to upgrade a mounted file system to disk layout Version 7 or later. Use the <code>vxfsconvert</code> command to upgrade an unmounted file system from disk layout Version 4 to disk layout Version 7 or later. Note: Symantec VirtualStore supports only disk layout Version 8 and 9.
Fast file system recovery	VxFS provides fast recovery of a file system from system failure. See “Fast file system recovery” on page 78.

Table 2-1 Veritas File System features in cluster file systems (*continued*)

Features	Description
FileSnap	A FileSnap is a space-optimized copy of a file in the same name space, stored in the same file system. See “About FileSnaps” on page 580.
Freeze and thaw	Synchronizing operations, which require freezing and thawing file systems, are done on a cluster-wide basis.
Locking	Advisory file and record locking are supported on SFCFSHA. For the <code>F_GETLK</code> command, if there is a process holding a conflicting lock, the <code>l_pid</code> field returns the process ID of the process holding the conflicting lock. The nodeid-to-node name translation can be done by examining the <code>/etc/llthosts</code> file or with the <code>fsclustadm</code> command. Mandatory locking, and deadlock detection supported by traditional <code>fcntl</code> locks, are not supported on SFCFSHA. See the <code>fcntl(2)</code> manual page.
Memory mapping	Shared memory mapping established by the <code>mmap()</code> function is supported on SFCFSHA. See the <code>mmap(2)</code> manual page.
NFS mounts	You export the NFS file systems from the Cluster. You can NFS export CFS file systems in a distributed highly available way.
Nested Mounts	You can use a directory on a cluster mounted file system as a mount point for a local file system or another cluster file system.
Quick I/O for databases	The Quick I/O for Databases feature, using clusterized Oracle Disk Manager (ODM), is supported on SFCFSHA. For DB2 and Sybase databases, see the <i>Veritas Storage Foundation Solution Guide</i>. For Oracle databases, see <i>Veritas Storage Foundation: Storage and Availability Management for Oracle Databases Guide</i>.
Quotas	Quotas are supported on cluster file systems.
Snapshots	Snapshots are supported on cluster file systems.
Storage Checkpoints	Storage Checkpoints are supported on cluster file systems, but are licensed only with other Veritas products.

Veritas File System features not in cluster file systems

Table 2-2 describes functionality as not supported and may not be expressly prevented from operating on cluster file systems, but the actual behavior is indeterminate.

It is not advisable to use unsupported functionality on SFCFSHA, or to alternate mounting file systems with these options as local and cluster mounts.

Table 2-2 Veritas File System features not in cluster file systems

Unsupported features	Comments
qllog	QuickLog is not supported.
Swap files	Swap files are not supported on cluster mounted file system.
mknod	The <code>mknod</code> command cannot be used to create devices on a cluster mounted file system.
Cache advisories	Cache advisories are set with the mount command on individual file systems, but are not propagated to other nodes of a cluster.
Cached Quick I/O	This Quick I/O for Databases feature that caches data in the file system cache is not supported.
Commands that depend on file access times	File access times may appear different across nodes because the <code>atime</code> file attribute is not closely synchronized in a cluster file system. So utilities that depend on checking access times may not function reliably.

About Storage Foundation Cluster File System High Availability and the Group Lock Manager

Storage Foundation Cluster File System High Availability (SFCFSHA) uses the Veritas Group Lock Manager (GLM) to reproduce UNIX single-host file system semantics in clusters. UNIX file systems make writes appear atomic. This means when an application writes a stream of data to a file, a subsequent application reading from the same area of the file retrieves the new data, even if it has been cached by the file system and not yet written to disk. Applications cannot retrieve stale data or partial results from a previous write.

To reproduce single-host write semantics, the file system must keep system caches coherent, and each must instantly reflect updates to cached data, regardless of the node from which the updates originate.

Storage Foundation Cluster File System High Availability namespace

The mount point name must remain the same for all nodes mounting the same cluster file system. This is required for the CFSSMount agent (online, offline, and monitoring) to work correctly.

About asymmetric mounts

A VxFS file system mounted with the `mount -o cluster` option is a cluster, or shared mount, as opposed to a non-shared or local mount. A file system mounted in shared mode must be on a VxVM shared volume in a cluster environment. A local mount cannot be remounted in shared mode and a shared mount cannot be remounted in local mode when you use the `mount -o remount` option. A single clustered file system can be mounted with different read/writes options on different nodes. These are called asymmetric mounts.

Asymmetric mounts allow shared file systems to be mounted with different read/write capabilities. For example, one node in the cluster can mount read/write, while other nodes mount read-only.

When a primary mounts "ro", this means that neither this node nor any other node is allowed to write the file system. Secondaries can only mount "ro", if the primary mounts "ro". Otherwise, the primary mounts either "rw" or "ro,crw", and the secondaries have the same choice.

You can specify the cluster read-write (`crw`) option when you first mount the file system, or the options can be altered when doing a remount (`mount -o remount`).

[Figure 2-1](#) describes the first column showing the mode in which the primary is mounted:

Figure 2-1 Primary and secondary mounts

		Secondary		
		ro	rw	ro, crw
Primary	ro	X		
	rw		X	X
	ro, crw		X	X

The check marks indicate the mode secondary mounts can use for a given mode of the primary.

Mounting the primary with only the `-o cluster,ro` option prevents the secondaries from mounting in a different mode; that is, read-write.

See the `mount_vxfs(1M)` manual page for more information.

Primary and secondary

A file system cluster consists of one primary, and up to 63 secondaries. The primary-secondary terminology applies to one file system, not to a specific node (or hardware platform). You can have the same cluster node be primary for one shared file system, while at the same time it is secondary for another shared file system. Such distribution of file system primaryship to balance the load on a cluster is a recommended administrative policy.

See [“Distribute the load on a cluster”](#) on page 362.

For CVM, a single cluster node is the master for all shared disk groups and shared volumes in the cluster.

Determining or moving primaryship

The first node of a cluster file system to mount is called the primary node. Other nodes are called secondary nodes. If a primary node fails, an internal election process determines which of the secondaries becomes the primary file system.

To determine primaryship

- To determine primaryship, type the following command:

```
# fsclustadm -v showprimary mount_point
```

To give primaryship to a node

- To give primaryship to a node, type the following command on the node:

```
# fsclustadm -v setprimary mount_point
```

Synchronize time on Cluster File Systems

SFCFSA requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for inode (`ctime`) and data modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

File system tunables

Using the `/etc/vx/tunefstab` file updates the tunable parameters at the time of mount. The file system `/etc/vx/tunefstab` parameters are set to be identical on all nodes by propagating the parameters to each cluster node. When the file system is mounted on the node, the `/etc/vx/tunefstab` parameters of the primary node are used. Symantec recommends that this file be identical on each node.

See the `tunefstab(4)` and `vxtunefs(1M)` manual pages for more information.

Setting the number of parallel fsck threads

This section describes how to set the number of parallel fsck threads.

The number of parallel fsck threads that could be active during recovery was set to 4. For example, if a node failed over 12 file systems, log replay for the 12 file systems will not complete at the same time. The number was set to 4 since parallel replay of a large number of file systems would put memory pressure on systems with less memory. However, on larger systems the restriction of 4 parallel processes replaying is not necessary.

This value gets tuned in accordance with available physical memory in the system.

To set the number of parallel fsck threads

- ◆ On all nodes in the cluster, edit the `/opt/VRTSvcs/bin/CFSfsckd/CFSfsckd.env` file and set `FCKD_OPTS="-n N"`.

where *N* is the number of parallel fsck threads desired and value of *N* has to be between 4 and 128.

Storage Checkpoints

Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) provides a Storage Checkpoint feature that quickly creates a persistent image of a file system at an exact point in time.

See [“About Storage Checkpoints”](#) on page 640.

Storage Foundation Cluster File System High Availability backup strategies

The same backup strategies used for standard VxFS can be used with SFCFSHA because the APIs and commands for accessing the namespace are the same. File System checkpoints provide an on-disk, point-in-time copy of the file system. Because performance characteristics of a checkpointed file system are better in certain I/O patterns, they are recommended over file system snapshots (described below) for obtaining a frozen image of the cluster file system.

File System snapshots are another method of a file system on-disk frozen image. The frozen image is non-persistent, in contrast to the checkpoint feature. A snapshot can be accessed as a read-only mounted file system to perform efficient online backups of the file system. Snapshots implement “copy-on-write” semantics that incrementally copy data blocks when they are overwritten on the snapped file system. Snapshots for cluster file systems extend the same copy-on-write mechanism for the I/O originating from any cluster node.

Mounting a snapshot filesystem for backups increases the load on the system because of the resources used to perform copy-on-writes and to read data blocks from the snapshot. In this situation, cluster snapshots can be used to do off-host backups. Off-host backups reduce the load of a backup application from the primary server. Overhead from remote snapshots is small when compared to overall snapshot overhead. Therefore, running a backup application by mounting a snapshot from a relatively less loaded node is beneficial to overall cluster performance.

The following are several characteristics of a cluster snapshot:

- A snapshot for a cluster mounted file system can be mounted on any node in a cluster. The file system can be a primary, secondary, or secondary-only. A stable image of the file system is provided for writes from any node.
See the `mount_vxfs` manual page for more information on secondary-only (seconly) is a CFS mount option.
- Multiple snapshots of a cluster file system can be mounted on the same or different cluster nodes.
- A snapshot is accessible only on the node mounting the snapshot. The snapshot device cannot be mounted on two nodes simultaneously.
- The device for mounting a snapshot can be a local disk or a shared volume. A shared volume is used exclusively by a snapshot mount and is not usable from other nodes as long as the snapshot is mounted on that device.
- On the node mounting a snapshot, the snapped file system cannot be unmounted while the snapshot is mounted.
- A SFCFSHA snapshot ceases to exist if it is unmounted or the node mounting the snapshot fails. However, a snapshot is not affected if another node leaves or joins the cluster.
- A snapshot of a read-only mounted file system cannot be taken. It is possible to mount a snapshot of a cluster file system only if the snapped cluster file system is mounted with the `crw` option.

In addition to frozen images of file systems, there are volume-level alternatives available for shared volumes using mirror split and rejoin. Features such as Fast Mirror Resync and Space Optimized snapshot are also available.

Parallel I/O

Some distributed applications read and write to the same file concurrently from one or more nodes in the cluster; for example, any distributed application where one thread appends to a file and there are one or more threads reading from various regions in the file. Several high-performance compute (HPC) applications can also benefit from this feature, where concurrent I/O is performed on the same file. Applications do not require any changes to use parallel I/O.

Traditionally, the entire file is locked to perform I/O to a small region. To support parallel I/O, SFCFSHA locks ranges in a file that correspond to I/O requests. The granularity of the locked range is a page. Two I/O requests conflict if at least one is a write request, and the I/O range of the request overlaps the I/O range of the other.

The parallel I/O feature enables I/O to a file by multiple threads concurrently, as long as the requests do not conflict. Threads issuing concurrent I/O requests could be executing on the same node, or on different nodes in the cluster.

An I/O request that requires allocation is not executed concurrently with other I/O requests. Note that when a writer is extending the file and readers are lagging behind, block allocation is not necessarily done for each extending write.

Predetermine the file size and preallocate the file to avoid block allocations during I/O. This improves the concurrency of applications performing parallel I/O to the file. Parallel I/O also avoids unnecessary page cache flushes and invalidations using range locking, without compromising the cache coherency across the cluster.

For applications that update the same file from multiple nodes, the `-nomtime` mount option provides further concurrency. Modification and change times of the file are not synchronized across the cluster, which eliminates the overhead of increased I/O and locking. The timestamp seen for these files from a node may not have the time updates that happened in the last 60 seconds.

I/O error handling policy

I/O errors can occur for several reasons, including failures of Fibre Channel links, host-bus adapters, and disks. SFCFSA disables the file system on the node encountering I/O errors. The file system remains available from other nodes.

After the hardware error is fixed (for example, the Fibre Channel link is reestablished), the file system can be force unmounted from all the active nodes in the cluster, and the mount resource can be brought online from the disabled node to reinstate the file system.

Recovering for I/O failures

The disabled file system can be restored by a force unmount and the resource will be brought online without rebooting, which also brings the shared disk group resource online.

Note: If the jeopardy condition is not fixed, the nodes are susceptible to leaving the cluster again on subsequent node failure.

See the *Veritas Cluster Server Administrator's Guide*.

Single network link and reliability

Certain environments may prefer using a single private link or a public network for connecting nodes in a cluster, despite the loss of redundancy for dealing with network failures. The benefits of this approach include simpler hardware topology and lower costs; however, there is obviously a tradeoff with high availability.

For the above environments, SFCFSA provides the option of a single private link, or using the public network as the private link if I/O fencing is present. I/O fencing is used to handle split-brain scenarios. The option for single network is given during installation.

See [“About preventing data corruption with I/O fencing”](#) on page 55.

Configuring a low-priority link

LLT can be configured to use a low-priority network link as a backup to normal heartbeat channels. Low-priority links are typically configured on a public or administrative network. This typically results in a completely different network infrastructure than the cluster private interconnect, and reduces the chance of a single point of failure bringing down all links. The low-priority link is not used for cluster membership traffic until it is the only remaining link. In normal operation, the low-priority link carries only heartbeat traffic for cluster membership and link state maintenance. The frequency of heartbeats drops 50 percent to reduce network overhead. When the low-priority link is the only remaining network link, LLT also switches over all cluster status traffic. Following repair of any configured private link, LLT returns cluster status traffic to the high-priority link.

LLT links can be added or removed while clients are connected. Shutting down GAB or the high-availability daemon, had, is not required.

To add a link

- To add a link, type the following command:

```
# lltconfig -d device -t device_tag
```

where *device_tag* is a tag to identify particular link in subsequent commands, and is displayed by `lltstat(1M)`.

To remove a link

- To remove a link, type the following command:

```
# lltconfig -u device_tag
```

See the `lltconfig(1M)` manual page.

Changes take effect immediately and are lost on the next reboot. For changes to span reboots you must also update the `/etc/llttab` file.

Note: LLT clients will not know how things are going until you only have one LLT link left and GAB declares jeopardy

Split-brain and jeopardy handling

A split-brain occurs when the cluster membership view differs among the cluster nodes, increasing the chance of data corruption. With I/O fencing, the potential for data corruption is eliminated. I/O fencing requires disks that support SCSI-3 PR.

You can also configure I/O fencing using coordination point servers (CP servers). In virtual environments that do not support SCSI-3, you can configure non-SCSI-3 server-based fencing.

See [“About server-based I/O fencing”](#) on page 66.

See [“About I/O fencing for SFCFSHA in virtual machines that do not support SCSI-3 PR”](#) on page 55.

Jeopardy state In the absence of I/O fencing, SFCFSHA installation requires two heartbeat links. When a node is down to a single heartbeat connection, SFCFSHA can no longer discriminate between loss of a system and loss of the final network connection. This state is defined as jeopardy.

SFCFSHA detects jeopardy and responds to it in ways that prevent data corruption in some split-brain situations. However, data corruption can still occur in other situations:

- All links go down simultaneously.
- A node hangs and is unable to respond to heartbeat messages.

To eliminate the chance of data corruption in these scenarios, I/O fencing is required. With I/O fencing, the jeopardy state does not require special handling by the SFCFSHA stack.

Jeopardy handling For installations that do not have I/O fencing configured, jeopardy handling prevents some potential split-brain conditions. If any cluster node fails following a jeopardy state notification, all cluster file systems that were mounted on the failed node or nodes are disabled on all remaining nodes. If a leave reconfiguration happens after jeopardy state notification, then the nodes which have received the jeopardy state notification leave the cluster.

About I/O fencing

I/O fencing protects the data on shared disks when nodes in a cluster detect a change in the cluster membership that indicates a split-brain condition.

The fencing operation determines the following:

- The nodes that must retain access to the shared storage
- The nodes that must be ejected from the cluster

This decision prevents possible data corruption. The installer installs the I/O fencing driver, VRTSvxfen, when you install SFCFSHA. To protect data on shared disks, you must configure I/O fencing after you install and configure SFCFSHA.

I/O fencing technology uses coordination points for arbitration in the event of a network partition.

I/O fencing coordination points can be coordinator disks or coordination point servers (CP servers) or both. You can configure disk-based or server-based I/O fencing:

Disk-based I/O fencing	I/O fencing that uses coordinator disks is referred to as disk-based I/O fencing. Disk-based I/O fencing ensures data integrity in a single cluster.
Server-based I/O fencing	I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing. Server-based fencing can include only CP servers, or a mix of CP servers and coordinator disks. Server-based I/O fencing ensures data integrity in multiple clusters. In virtualized environments that do not support SCSI-3 PR, SFCFSHA supports non-SCSI-3 server-based I/O fencing.

See [“About preventing data corruption with I/O fencing”](#) on page 55.

Note: Symantec recommends that you use I/O fencing to protect your cluster against split-brain situations.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

About I/O fencing for SFCFSHA in virtual machines that do not support SCSI-3 PR

In a traditional I/O fencing implementation, where the coordination points are coordination point servers (CP servers) or coordinator disks, Veritas Clustered Volume Manager and Veritas I/O fencing modules provide SCSI-3 persistent reservation (SCSI-3 PR) based protection on the data disks. This SCSI-3 PR protection ensures that the I/O operations from the losing node cannot reach a disk that the surviving sub-cluster has already taken over.

In virtualized environments that do not support SCSI-3 PR, SFCFSHA attempts to provide reasonable safety for the data disks. SFCFSHA requires you to configure non-SCSI-3 server-based I/O fencing in such environments. Non-SCSI-3 fencing uses CP servers as coordination points with some additional configuration changes to support I/O fencing in such environments.

About preventing data corruption with I/O fencing

I/O fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster.

To provide high availability, the cluster must be capable of taking corrective action when a node fails. In this situation, SFCFSHA configures its components to reflect the altered membership.

Problems arise when the mechanism that detects the failure breaks down because symptoms appear identical to those of a failed node. For example, if a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner.

In addition to a broken set of private networks, other scenarios can generate this situation. If a system is so busy that it appears to stop responding or "hang," the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

SFCFSHA uses I/O fencing to remove the risk that is associated with split-brain. I/O fencing allows write access for members of the active cluster. It blocks access to storage from non-members.

About SCSI-3 Persistent Reservations

SCSI-3 Persistent Reservations (SCSI-3 PR) are required for I/O fencing and resolve the issues of using SCSI reservations in a clustered SAN environment. SCSI-3 PR enables access for multiple nodes to a device and simultaneously blocks access for other nodes.

SCSI-3 reservations are persistent across SCSI bus resets and support multiple paths from a host to a disk. In contrast, only one host can use SCSI-2 reservations with one path. If the need arises to block access to a device because of data integrity concerns, only one host and one path remain active. The requirements for larger clusters, with multiple nodes reading and writing to storage in a controlled manner, make SCSI-2 reservations obsolete.

SCSI-3 PR uses a concept of registration and reservation. Each system registers its own "key" with a SCSI-3 device. Multiple systems registering keys form a membership and establish a reservation, typically set to "Write Exclusive Registrants Only." The WERO setting enables only registered systems to perform write operations. For a given disk, only one reservation can exist amidst numerous registrations.

With SCSI-3 PR technology, blocking write access is as easy as removing a registration from a device. Only registered members can "eject" the registration of another member. A member wishing to eject another member issues a "preempt and abort" command. Ejecting a node is final and atomic; an ejected node cannot eject another node. In SFCFSHA, a node registers the same key for all paths to the device. A single preempt and abort command ejects a node from all paths to the storage device.

About I/O fencing operations

I/O fencing, provided by the kernel-based fencing module (vxfen), performs identically on node failures and communications failures. When the fencing module on a node is informed of a change in cluster membership by the GAB module, it immediately begins the fencing operation. The node tries to eject the key for departed nodes from the coordinator disks using the preempt and abort command. When the node successfully ejects the departed nodes from the coordinator disks, it also ejects the departed nodes from the data disks. In a split-brain scenario, both sides of the split would race for control of the coordinator disks. The side winning the majority of the coordinator disks wins the race and fences the loser. The loser then panics and restarts the system.

About I/O fencing components

The shared storage for SFCFSHA must support SCSI-3 persistent reservations to enable I/O fencing. SFCFSHA involves two types of shared storage:

- Data disks—Store shared data
See “[About data disks](#)” on page 57.
- Coordination points—Act as a global lock during membership changes
See “[About coordination points](#)” on page 57.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs).

These disks must support SCSI-3 PR and must be part of standard VxVM or CVM disk groups. CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. SFCFSA prevents split-brain when vxfen races for control of the coordination points and the winner partition fences the ejected nodes from accessing the data disks.

Note: Typically, a fencing configuration for a cluster must have three coordination points. Symantec also supports server-based fencing with a single CP server as its only coordination point with a caveat that this CP server becomes a single point of failure.

The coordination points can either be disks or servers or both.

- Coordinator disks
Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the SFCFSA configuration.
You can configure coordinator disks to use Veritas Volume Manager Dynamic Multi-pathing (DMP) feature. Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use either DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is dmp by default.
See the *Veritas Storage Foundation Administrator's Guide*.

- Coordination point servers

The coordination point server (CP server) is a software solution which runs on a remote system or cluster. CP server provides arbitration functionality by allowing the SFCFSHA cluster nodes to perform the following tasks:

- Self-register to become a member of an active SFCFSHA cluster (registered with CP server) with access to the data drives
- Check which other nodes are registered as members of this active SFCFSHA cluster
- Self-unregister from this active SFCFSHA cluster
- Forcefully unregister other nodes (preempt) as members of this active SFCFSHA cluster

In short, the CP server functions as another arbitration mechanism that integrates within the existing I/O fencing module.

Note: With the CP server, the fencing arbitration logic still remains on the SFCFSHA cluster.

Multiple SFCFSHA clusters running different operating systems can simultaneously access the CP server. TCP/IP based communication is used between the CP server and the SFCFSHA clusters.

About preferred fencing

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points. With the preferred fencing feature, you can specify how the fencing driver must determine the surviving subcluster.

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy as follows:

- Enable system-based preferred fencing policy to give preference to high capacity systems.
- Enable group-based preferred fencing policy to give preference to service groups for high priority applications.
- Disable preferred fencing policy to use the default node count-based race policy.

See [“How preferred fencing works”](#) on page 59.

See [“Enabling or disabling the preferred fencing policy”](#) on page 453.

How preferred fencing works

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. At the time of a network partition, the fencing driver in each subcluster races for the coordination points. The subcluster that grabs the majority of coordination points survives whereas the fencing driver causes a system panic on nodes from all other subclusters. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points.

This default racing preference does not take into account the application groups that are online on any nodes or the system capacity in any subcluster. For example, consider a two-node cluster where you configured an application on one node and the other node is a standby-node. If there is a network partition and the standby-node wins the race, the node where the application runs panics and VCS has to bring the application online on the standby-node. This behavior causes disruption and takes time for the application to fail over to the surviving node and then to start up again.

The preferred fencing feature lets you specify how the fencing driver must determine the surviving subcluster. The preferred fencing solution makes use of a fencing parameter called node weight. VCS calculates the node weight based on online applications and system capacity details that you provide using specific VCS attributes, and passes to the fencing driver to influence the result of race for coordination points. At the time of a race, the racer node adds up the weights for all nodes in the local subcluster and in the leaving subcluster. If the leaving subcluster has a higher sum (of node weights) then the racer for this subcluster delays the race for the coordination points. Thus, the subcluster that has critical systems or critical applications wins the race.

The preferred fencing feature uses the cluster-level attribute PreferredFencingPolicy that takes the following race policy values:

- **Disabled (default):** Preferred fencing is disabled.
When the PreferredFencingPolicy attribute value is set as Disabled, VCS sets the count based race policy and resets the value of node weight as 0.
- **System:** Based on the capacity of the systems in a subcluster.
If one system is more powerful than others in terms of architecture, number of CPUs, or memory, this system is given preference in the fencing race.
When the PreferredFencingPolicy attribute value is set as System, VCS calculates node weight based on the system-level attribute FencingWeight.
- **Group:** Based on the higher priority applications in a subcluster.
The fencing driver takes into account the service groups that are online on the nodes in any subcluster. In the event of a network partition, the node with higher priority service groups is given preference in the fencing race.

When the PreferredFencingPolicy attribute value is set as Group, VCS calculates node weight based on the group-level attribute Priority for those service groups that are active.

See the *Veritas Cluster Server Administrator's Guide* for more information on the VCS attributes.

See [“Enabling or disabling the preferred fencing policy”](#) on page 453.

About I/O fencing configuration files

[Table 2-3](#) lists the I/O fencing configuration files.

Table 2-3 I/O fencing configuration files

File	Description
/etc/default/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFCFSHA configuration.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing.

Table 2-3 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	This file contains the following parameters: ■ vxfen_mode ■ scsi3—For disk-based fencing ■ customized—For server-based fencing ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ vxfen_mechanism This parameter is applicable only for server-based fencing. Set the value as cps. ■ scsi3_disk_policy ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. ■ raw—Configure the vxfen module to use the underlying raw character devices Note: You must use the same SCSI-3 disk policy on all the nodes. ■ security This parameter is applicable only for server-based fencing. 1—Indicates that communication with the CP server is in secure mode. This setting is the default. 0—Indicates that communication with the CP server is in non-secure mode. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in a server-based fencing can include coordinator disks, CP servers, or a mix of both. If you use coordinator disks, you must create a coordinator disk group with the coordinator disk names. Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ single_cp This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk. ■ autoseed_gab_timeout This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature requires that I/O fencing is enabled. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table 2-3 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points. Note: The /etc/vxfentab file is a generated file; do not modify this file. For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows: - Raw disk: <pre> /dev/rdisk/c1t1d0s2 /dev/rdisk/c2t1d0s2 /dev/rdisk/c3t1d2s2 </pre> - DMP disk: <pre> /dev/vx/rdmp/c1t1d0s2 /dev/vx/rdmp/c2t1d0s2 /dev/vx/rdmp/c3t1d0s2 </pre> For server-based fencing, the /etc/vxfentab file also includes the security settings information. For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.

How I/O fencing works in different event scenarios

[Table 2-4](#) describes how I/O fencing works to prevent data corruption in different failure event scenarios. For each event, review the corrective operator actions.

Table 2-4 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Both private networks fail.	Node A races for majority of coordination points. If Node A wins race for coordination points, Node A ejects Node B from the shared disks and continues.	Node B races for majority of coordination points. If Node B loses the race for the coordination points, Node B panics and removes itself from the cluster.	When Node B is ejected from cluster, repair the private networks before attempting to bring Node B back.
Both private networks function again after event above.	Node A continues to work.	Node B has crashed. It cannot start the database since it is unable to write to the data disks.	Restart Node B after private networks are restored.
One private network fails.	Node A prints message about an IOFENCE on the console but continues.	Node B prints message about an IOFENCE on the console but continues.	Repair private network. After network is repaired, both nodes automatically use it.
Node A hangs.	Node A is extremely busy for some reason or is in the kernel debugger. When Node A is no longer hung or in the kernel debugger, any queued writes to the data disks fail because Node A is ejected. When Node A receives message from GAB about being ejected, it panics and removes itself from the cluster.	Node B loses heartbeats with Node A, and races for a majority of coordination points. Node B wins race for coordination points and ejects Node A from shared data disks.	Repair or debug the node that hangs and reboot the node to rejoin the cluster.

Table 2-4 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
Nodes A and B and private networks lose power. Coordination points and data disks retain power. Power returns to nodes and they restart, but private networks still have no power.	Node A restarts and I/O fencing driver (vxfs) detects Node B is registered with coordination points. The driver does not see Node B listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node A from joining the cluster. Node A console displays: Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.	Node B restarts and I/O fencing driver (vxfs) detects Node A is registered with coordination points. The driver does not see Node A listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node B from joining the cluster. Node B console displays: Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.	Resolve preexisting split-brain condition.

Table 2-4 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
Node A crashes while Node B is down. Node B comes up and Node A is still down.	Node A is crashed.	Node B restarts and detects Node A is registered with the coordination points. The driver does not see Node A listed as member of the cluster. The I/O fencing device driver prints message on console: Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.	Resolve preexisting split-brain condition.
The disk array containing two of the three coordination points is powered off. No node leaves the cluster membership	Node A continues to operate as long as no nodes leave the cluster.	Node B continues to operate as long as no nodes leave the cluster.	Power on the failed disk array so that subsequent network partition does not cause cluster shutdown, or replace coordination points. See “Replacing I/O fencing coordinator disks when the cluster is online” on page 408.

Table 2-4 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
The disk array containing two of the three coordination points is powered off. Node B gracefully leaves the cluster and the disk array is still powered off. Leaving gracefully implies a clean shutdown so that vxfen is properly unconfigured.	Node A continues to operate in the cluster.	Node B has left the cluster.	Power on the failed disk array so that subsequent network partition does not cause cluster shutdown, or replace coordination points. See “Replacing I/O fencing coordinator disks when the cluster is online” on page 408.
The disk array containing two of the three coordination points is powered off. Node B abruptly crashes or a network partition occurs between node A and node B, and the disk array is still powered off.	Node A races for a majority of coordination points. Node A fails because only one of the three coordination points is available. Node A panics and removes itself from the cluster.	Node B has left cluster due to crash or network partition.	Power on the failed disk array and restart I/O fencing driver to enable Node A to register with all coordination points, or replace coordination points.

About server-based I/O fencing


In a disk-based I/O fencing implementation, the vxfen driver handles various SCSI-3 PR based arbitration operations completely within the driver. I/O fencing also provides a framework referred to as customized fencing wherein arbitration operations are implemented in custom scripts. The vxfen driver invokes the custom scripts.

The CP server-based coordination point uses a customized fencing framework. Note that SCSI-3 PR based fencing arbitration can also be enabled using customized fencing framework. This allows the user to specify a combination of SCSI-3 LUNs and CP servers as coordination points using customized fencing. Customized fencing can be enabled by specifying vxfen_mode=customized and vxfen_mechanism=cps in the `/etc/vxfenmode` file.

Moreover, both `/etc/vxfenmode` and `/etc/vxfentab` files contain additional parameter "security" which indicates if communication between CP server and SFCFSHA cluster nodes is secure.

Figure 2-2 displays a schematic of the customized fencing options.

Figure 2-2 Customized fencing



A user level daemon `vxfend` interacts with the `vxfen` driver, which in turn interacts with `GAB` to get the node membership update. Upon receiving membership updates, `vxfend` invokes various scripts to race for the coordination point and fence off data disks. The `vxfend` daemon manages various fencing agents. The customized fencing scripts are located in the `/opt/VRTSvcs/vxfen/bin/customized/cps` directory.

The scripts that are involved include the following:

- `generate_snapshot.sh` : Retrieves the SCSI ID's of the coordinator disks and/or UUID ID's of the CP servers

CP server uses the UUID stored in `/etc/VRTScps/db/current/cps_uuid`.

For information about the UUID (Universally Unique Identifier), see the *Veritas Cluster Server Administrator's Guide*.

- `join_local_node.sh`: Registers the keys with the coordinator disks or CP servers
- `race_for_coordination_point.sh`: Races to determine a winner after cluster reconfiguration
- `unjoin_local_node.sh`: Removes the keys that are registered in `join_local_node.sh`
- `fence_data_disks.sh`: Fences the data disks from access by the losing nodes.
- `local_info.sh`: Lists local node's configuration parameters and coordination points, which are used by the vxfen driver.

I/O fencing enhancements provided by CP server


CP server configurations enhance disk-based I/O fencing by providing the following new capabilities:

- CP server configurations are scalable, and a configuration with three CP servers can provide I/O fencing for multiple SFCFSHA clusters. Since a single CP server configuration can serve a large number of SFCFSHA clusters, the cost of multiple SFCFSHA cluster deployments can be significantly reduced.
- Appropriately situated CP servers can eliminate any coordinator disk location bias in the I/O fencing process. For example, this location bias may occur where, due to logistical restrictions, two of the three coordinator disks are located at a single site, and the cost of setting up a third coordinator disk location is prohibitive.

See [Figure 2-3](#) on page 69.

In such a configuration, if the site with two coordinator disks is inaccessible, the other site does not survive due to a lack of a majority of coordination points. I/O fencing would require extension of the SAN to the third site which may not be a suitable solution. An alternative is to place a CP server at a remote site as the third coordination point.

Note: The CP server provides an alternative arbitration mechanism without having to depend on SCSI-3 compliant coordinator disks. Data disk fencing in CVM will still require SCSI-3 I/O fencing.

Figure 2-3 Skewed placement of coordinator disks at Site 1

About the CP server database

CP server requires a database for storing the registration keys of the SFCFSA cluster nodes. CP server uses a SQLite database for its operations. By default, the database is located at `/etc/VRTScps/db`.

For a single node VCS cluster hosting a CP server, the database can be placed on a local file system. For an SFHA cluster hosting a CP server, the database must be placed on a shared file system. The file system must be shared among all nodes that are part of the SFHA cluster.

In an SFHA cluster hosting the CP server, the shared database is protected by setting up SCSI-3 PR based I/O fencing. SCSI-3 PR based I/O fencing protects against split-brain scenarios.

Warning: The CP server database must not be edited directly and should only be accessed using `cpsadm(1M)`. Manipulating the database manually may lead to undesirable results including system panics.

About the CP server user types and privileges

The CP server supports the following user types, each with a different access level privilege:

- CP server administrator (admin)
- CP server operator

Different access level privileges permit the user to issue different commands. If a user is neither a CP server admin nor a CP server operator user, then the user has guest status and can issue limited commands.

The user types and their access level privileges are assigned to individual users during SFCFSHA cluster configuration for fencing. During the installation process, you are prompted for a user name, password, and access level privilege (CP server admin or CP server operator).

To administer and operate a CP server, there must be at least one CP server admin.

A root user on a CP server is given all the administrator privileges, and these administrator privileges can be used to perform all the CP server specific operations.

About secure communication between the SFCFSHA cluster and CP server

In a data center, TCP/IP communication between the SFCFSHA cluster (application cluster) and CP server must be made secure. The security of the communication channel involves encryption, authentication, and authorization.

The CP server node or cluster needs to confirm the authenticity of the SFCFSHA cluster nodes that communicate with it as a coordination point and only accept requests from known SFCFSHA cluster nodes. Requests from unknown clients are rejected as non-authenticated. Similarly, the fencing framework in SFCFSHA cluster must confirm that authentic users are conducting fencing operations with the CP server.

Entities on behalf of which authentication is done, are referred to as principals. On the SFCFSHA cluster nodes, the current VCS installer creates the Authentication Server credentials on each node in the cluster. It also creates vcsauthserver which authenticates the credentials. The installer then proceeds to start VCS in secure mode.

Typically, in an existing VCS cluster with security configured, vcsauthserver runs on each cluster node.


How secure communication between the CP servers and the SFCFSHA clusters work

CP server and SFCFSHA cluster (application cluster) node communication involve the following entities:

- vxcpserv for the CP server
- cpsadm for the SFCFSHA cluster node

Figure 2-4 displays a schematic of the end-to-end communication flow with security enabled on CP server and SFCFSHA clusters (application clusters).

Figure 2-4 End-To-end communication flow with security enabled on CP server and SFCFSHA clusters



Communication flow between CP server and SFCFSHA cluster nodes with security configured on them is as follows:

- Initial setup:
Identities of CP server and SFCFSHA cluster nodes are configured on respective nodes by the VCS installer.

Note: At the time of fencing configuration, the installer establishes trust between the CP server and the application cluster so that vxcperv process can authenticate requests from the application cluster nodes. If you manually configured I/O fencing, then you must set up trust between the CP server and the application cluster.

The `cpsadm` command gets the user name, domain type from the environment variables `CPS_USERNAME`, `CPS_DOMAINTYPE`. The user is expected to export these variables before running the `cpsadm` command manually. The customized fencing framework exports these environment variables internally before running the `cpsadm` commands.

The CP server process (vxcperv) uses its own user (CPSEVER) which is added to the local vcsauthserver.

- Getting credentials from authentication broker:

The `cpsadm` command tries to get the existing credentials that is present on the local node. The installer generates these credentials during fencing configuration. The `vxcpserv` process tries to get the existing credentials that is present on the local node. The installer generates these credentials when it enables security.

- Communication between CP server and SFCFSHA cluster nodes:
After the CP server establishes its credential and is up, it becomes ready to receive data from the clients. After the `cpsadm` command obtains its credentials and authenticates CP server credentials, `cpsadm` connects to the CP server. Data is passed over to the CP server.
- Validation:
On receiving data from a particular SFCFSHA cluster node, `vxcpserv` validates its credentials. If validation fails, then the connection request data is rejected.

Security configuration details on CP server and SFCFSHA cluster

This section discusses the security configuration details for the CP server and SFCFSHA cluster (application cluster).

Settings in secure mode

The following are the settings for secure communication between the CP server and SFCFSHA cluster:

- CP server settings:
Installer creates a user with the following values:
 - username: CPSERVER
 - domainname: VCS_SERVICES@*cluster_uuid*
 - domaintype: vx

Run the following commands on the CP server to verify the settings:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSERVER  
  
# /opt/VRTScps/bin/cpsat showcred
```

Note: The CP server configuration file (`/etc/vxcps.conf`) must not contain a line specifying **security=0**. If there is no line specifying "security" parameter or if there is a line specifying **security=1**, CP server with security is enabled (which is the default).

- SFCFSHA cluster node(s) settings:

On SFCFSHA cluster, the installer creates a user for cpsadm during fencing configuration with the following values:

- username: CPSADM
- domainname: VCS_SERVICES@*cluster_uuid*
- domaintype: vx

Run the following commands on the SFCFSHA cluster node(s) to verify the security settings:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSADM
# /opt/VRTSvcs/bin/cpsat showcred
```

The users described above are used only for authentication for the communication between the CP server and the SFCFSHA cluster nodes.

For CP server's authorization, customized fencing framework on the SFCFSHA cluster uses the following user if security is configured:

CPSADM@VCS_SERVICES@*cluster_uuid*

where *cluster_uuid* is the application cluster's universal unique identifier.

For each SFCFSHA cluster node, this user must be registered on the CP server database before fencing starts on the SFCFSHA cluster node(s). This can be verified by issuing the following command:

```
# cpsadm -s cp_server -a list_users
```

The following is an example of the command output:

```
Username/Domain Type
CPSADM@VCS_SERVICES@77a2549c-1dd2-11b2-88d6-00306e4b2e0b/vx

Cluster Name / UUID                               Role
cluster1/{77a2549c-1dd2-11b2-88d6-00306e4b2e0b} Operator
```

Note: The configuration file (/etc/vxfenmode) on each client node must not contain a line specifying **security=0**. If there is no line specifying "security" parameter or if there is a line specifying **security=1**, client node starts with security enabled (which is the default).

Settings in non-secure mode

In non-secure mode, only authorization is provided on the CP server. Passwords are not requested. Authentication and encryption are not provided. User credentials

of “cpsclient@hostname” of “vx” domain type are used by the customized fencing framework for communication between CP server or SFCFSHA cluster node(s).

For each SFCFSHA cluster node, this user must be added on the CP server database before fencing starts on the SFCFSHA cluster node(s). The user can be verified by issuing the following command:

```
# cpsadm -s cpserver -a list_users
```

The following is an example of the command output:

Username/Domain Type	Cluster Name / UUID	Role
cpsclient@galaxy/vx	cluster1 / {f0735332-e3709c1c73b9}	Operator

Note: In non-secure mode, CP server configuration file (/etc/vxcps.conf) should contain a line specifying **security=0**. Similarly, on each SFCFSHA cluster node the configuration file (/etc/vxfenmode) should contain a line specifying **security=0**.

Storage Foundation Cluster File System High Availability and Veritas Volume Manager cluster functionality agents

Agents are VCS processes that manage predefined resource types. SFCFSHA and CVM require agents to interact with VCS. Agents bring resources online, take resources offline, monitor resources, and report any state changes to VCS. VCS bundled agents are part of VCS and are installed when VCS is installed. The SFCFSHA and CVM agents are add-on resources to VCS specifically for the Veritas File System and Veritas Volume Manager.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

Veritas Volume Manager cluster functionality

The Veritas Volume Manager cluster functionality (CVM) makes logical volumes accessible throughout a cluster. CVM enables multiple hosts to concurrently access the logical volumes under its control. A VxVM cluster comprises nodes sharing a set of devices. The nodes are connected across a network. If one node fails, other nodes can access the devices. The VxVM cluster feature presents the same logical view of the device configurations, including changes, on all nodes. You configure CVM shared storage after VCS sets up a cluster configuration.

How Veritas File System works

This chapter includes the following topics:

- [Veritas File System features](#)
- [Veritas File System performance enhancements](#)
- [Using Veritas File System](#)

Veritas File System features

VxFS includes the following features:

- [Extent-based allocation](#)
Extents allow disk I/O to take place in units of multiple blocks if storage is allocated in contiguous blocks.
- [Extent attributes](#)
Extent attributes are the extent allocation policies associated with a file.
- [Fast file system recovery](#)
VxFS provides fast recovery of a file system from system failure.
- [Extended mount options](#)
The VxFS file system supports extended `mount` options to specify enhanced data integrity modes, enhanced performance modes, temporary file system modes, improved synchronous writes, and large file sizes.
- [Enhanced data integrity modes](#)
VxFS avoids the problem of uninitialized data appearing in a file by waiting until the data has been flushed to disk before updating the new file size to disk.
- [Enhanced performance mode](#)

VxFS provides mount options to improve performance.

- **Temporary file system mode**
VxFS supplies an option to allow users to achieve higher performance on temporary file systems by delaying the logging for most operations.
- **Improved synchronous writes**
VxFS provides superior performance for synchronous write applications.
- **Support for large files and large file systems**
VxFS supports files larger than two gigabytes and large file systems up to 256 terabytes.
- **Access Control Lists**
An Access Control List (ACL) stores a series of entries that identify specific users or groups and their access privileges for a directory or file.
- **Storage Checkpoints**
Backup and restore applications can leverage Storage Checkpoints, a disk- and I/O-efficient copying technology for creating periodic frozen images of a file system.
- **FileSnaps**
A FileSnap is a space-optimized copy of a file in the same name space, stored in the same file system.
- **Online backup**
VxFS provides online data backup using the snapshot feature.
- **Quotas**
VxFS supports quotas, which allocate per-user and per-group quotas and limit the use of two principal resources: files and data blocks.
- **Cluster file systems**
Clustered file systems are an extension of VxFS that support concurrent direct media access from multiple systems.
- **Support for databases**
Databases can be created on files exported as character devices to achieve the same performance as databases created on raw disks.
- **Cross-platform data sharing**
Cross-platform data sharing allows data to be serially shared among heterogeneous systems where each system has direct access to the physical devices that hold the data.
- **File Change Log**
The VxFS File Change Log tracks changes to files and directories in a file system.

- [Reverse path name lookup](#)
The reverse path name lookup feature obtains the full path name of a file or directory from the inode number of that file or directory.
- [Multi-volume file system support](#)
The multi-volume support feature allows several volumes to be represented by a single logical object. VxFS supports creating file system on such multi-volumes.
- [SmartTier](#)
The SmartTier option allows you to configure policies that automatically relocate files from one volume to another, or relocate files by running file relocation commands, which can improve performance for applications that access specific types of files.

Note: In the previous VxFS 5.x releases, SmartTier was known as Dynamic Storage Tiering.

- [Thin Reclamation of a file system](#)
The Thin Reclamation feature allows you to release free data blocks of a VxFS file system to the free storage pool of a Thin Storage LUN. This feature is only supported on file systems created on a VxVM volume.
- [Partitioned directories](#)
The partitioned directories feature improves the directory performance of file systems. This feature operates only on disk layout Version 8 or later file systems.
- [Data deduplication on Solaris SPARC](#)
You can perform post-process periodic deduplication in a file system to eliminate duplicate data without any continuous cost.
- [File compression](#)
Compressing files reduces the space used by files, while retaining the accessibility of the files and being transparent to applications.

Note: VxFS does not support linking, removing, or renaming "." and ".." directory entries. These operations may disrupt file system operations.

Extent-based allocation

Disk space is allocated in 512-byte sectors to form logical blocks. VxFS supports logical block sizes of 1024, 2048, 4096, and 8192 bytes. The default block size is 1K for file system sizes of up to 1 TB, and 8K for file system sizes 1 TB or larger.

An extent is defined as one or more adjacent blocks of data within the file system. An extent is presented as an address-length pair, which identifies the starting block address and the length of the extent (in file system or logical blocks). VxFS allocates storage in extents rather than a block at a time.

Extents allow disk I/O to take place in units of multiple blocks if storage is allocated in contiguous blocks. For sequential I/O, multiple block operations are considerably faster than block-at-a-time operations; almost all disk drives accept I/O operations on multiple blocks.

Extent allocation only slightly alters the interpretation of addressed blocks from the inode structure compared to block based inodes. A VxFS inode references 10 direct extents, each of which are pairs of starting block addresses and lengths in blocks.

Extent attributes

VxFS allocates disk space to files in groups of one or more extents. VxFS also allows applications to control some aspects of the extent allocation. Extent attributes are the extent allocation policies associated with a file.

The `setext` and `getext` commands allow the administrator to set or view extent attributes associated with a file, as well as to preallocate space for a file.

Because these facilities are provided using VxFS-specific IOCTL system calls, most existing UNIX system applications do not use them. The VxFS-specific `cp`, `cpio`, and `mv` utilities use the facilities to preserve extent attributes and allocate space more efficiently. The current attributes of a file can be listed using the `getext` or VxFS-specific `ls` command. The facilities can also improve performance for custom applications. For portability reasons, these applications must check which file system type they are using before using these interfaces.

See the `setext(1)` and `getext(1)` manual pages.

Fast file system recovery

Most file systems rely on full structural verification by the `fsck` utility as the only means to recover from a system failure. For large disk configurations, this involves a time-consuming process of checking the entire structure, verifying that the file system is intact, and correcting any inconsistencies. VxFS provides fast recovery with the VxFS intent log and VxFS intent log resizing features.

VxFS intent log

VxFS reduces system failure recovery times by tracking file system activity in the VxFS intent log. This feature records pending changes to the file system structure in a circular intent log. The intent log recovery feature is not readily apparent to

users or a system administrator except during a system failure. During system failure recovery, the VxFS `fsck` utility performs an intent log replay, which scans the intent log and nullifies or completes file system operations that were active when the system failed. The file system can then be mounted without requiring a full structural check of the entire file system. Replaying the intent log may not completely recover the damaged file system structure if there was a disk hardware failure; hardware problems may require a complete system check using the `fsck` utility provided with VxFS.

The `mount` command automatically runs the VxFS `fsck` command to perform an intent log replay if the `mount` command detects a dirty log in the file system. This functionality is only supported on a file system mounted on a Veritas Volume Manager (VxVM) volume, and is supported on cluster file systems.

See [“The log option and data integrity”](#) on page 80.

VxFS intent log resizing

The VxFS intent log is allocated when the file system is first created. The size of the intent log is based on the size of the file system—the larger the file system, the larger the intent log. The maximum default intent log size for disk layout Version 7 or later is 256 megabytes.

Note: Inappropriate sizing of the intent log can have a negative impact on system performance.

See the `mkfs_vxfs(1M)` and the `fsadm_vxfs(1M)` manual pages.

Extended mount options

The VxFS file system provides the following enhancements to the `mount` command:

- Enhanced data integrity modes
- Enhanced performance mode
- Temporary file system mode
- Improved synchronous writes
- Support for large file sizes

See [“Mounting a VxFS file system”](#) on page 215.

See the `mount_vxfs(1M)` manual page.

Enhanced data integrity modes

For most UNIX file systems, including VxFS, the default mode for writing to a file is delayed, or buffered, meaning that the data to be written is copied to the file system cache and later flushed to disk.

A delayed write provides much better performance than synchronously writing the data to disk. However, in the event of a system failure, data written shortly before the failure may be lost since it was not flushed to disk. In addition, if space was allocated to the file as part of the write request, and the corresponding data was not flushed to disk before the system failure occurred, uninitialized data can appear in the file.

For the most common type of write, delayed extending writes (a delayed write that increases the file size), VxFS avoids the problem of uninitialized data appearing in the file by waiting until the data has been flushed to disk before updating the new file size to disk. If a system failure occurs before the data has been flushed to disk, the file size has not yet been updated, thus no uninitialized data appears in the file. The unused blocks that were allocated are reclaimed.

The `blkclear` option and data integrity

In environments where performance is more important than absolute data integrity, the preceding situation is not of great concern. However, VxFS supports environments that emphasize data integrity by providing the `mount -o blkclear` option that ensures uninitialized data does not appear in a file.

The `closesync` option and data integrity

VxFS provides the `mount -o mincache=closesync` option, which is useful in desktop environments with users who are likely to shut off the power on machines without halting them first. In `closesync` mode, only files that are written during the system crash or shutdown can lose data. Any changes to a file are flushed to disk when the file is closed.

The `log` option and data integrity

File systems are typically asynchronous in that structural changes to the file system are not immediately written to disk, which provides better performance. However, recent changes made to a system can be lost if a system failure occurs. Specifically, attribute changes to files and recently created files may disappear.

The `mount -o log` intent logging option guarantees that all structural changes to the file system are logged to disk before the system call returns to the application. With this option, the `rename(2)` system call flushes the source file to disk to guarantee the persistence of the file data before renaming it. The `rename()` call is also

guaranteed to be persistent when the system call returns. The changes to file system data and metadata caused by the `fsync(2)` and `fdatasync(2)` system calls are guaranteed to be persistent once the calls return.

Enhanced performance mode

VxFS has a mount option that improves performance: `delaylog`.

The `delaylog` option and enhanced performance

The default VxFS logging mode, `mount -o delaylog`, increases performance by delaying the logging of some structural changes. However, `delaylog` does not provide the equivalent data integrity as the previously described modes because recent changes may be lost during a system failure. This option provides at least the same level of data accuracy that traditional UNIX file systems provide for system failures, along with fast file system recovery.

Temporary file system mode

On most UNIX systems, temporary file system directories, such as `/tmp` and `/usr/tmp`, often hold files that do not need to be retained when the system reboots. The underlying file system does not need to maintain a high degree of structural integrity for these temporary directories. VxFS provides the `mount -o tmplog` option, which allows the user to achieve higher performance on temporary file systems by delaying the logging of most operations.

Improved synchronous writes

VxFS provides superior performance for synchronous write applications. The `mount -o datainlog` option greatly improves the performance of small synchronous writes.

The `mount -o convosync=dsync` option improves the performance of applications that require synchronous data writes but not synchronous inode time updates.

Warning: The use of the `-o convosync=dsync` option violates POSIX semantics.

Support for large files and large file systems

With VxFS, you can create, mount, and manage file systems containing large files (files larger than two gigabytes).

Warning: Some applications and utilities may not work on large files.

Access Control Lists

An Access Control List (ACL) stores a series of entries that identify specific users or groups and their access privileges for a directory or file. A file may have its own ACL or may share an ACL with other files. ACLs have the advantage of specifying detailed access permissions for multiple users and groups. On Solaris SPARC, ACLs are supported on cluster file systems.

See the `getfacl(1)` and `setfacl(1)` manual pages.

Storage Checkpoints

To increase availability, recoverability, and performance, Veritas File System offers on-disk and online backup and restore capabilities that facilitate frequent and efficient backup strategies. Backup and restore applications can leverage a Storage Checkpoint, a disk- and I/O-efficient copying technology for creating periodic frozen images of a file system. Storage Checkpoints present a view of a file system at a point in time, and subsequently identifies and maintains copies of the original file system blocks. Instead of using a disk-based mirroring method, Storage Checkpoints save disk space and significantly reduce I/O overhead by using the free space pool available to a file system.

Storage Checkpoint functionality is separately licensed.

See [“About Storage Checkpoints”](#) on page 640.

FileSnaps

A FileSnap is a space-optimized copy of a file in the same name space, stored in the same file system. Veritas File System (VxFS) supports FileSnaps on file systems with disk layout Version 8 or later.

See the `vxfilesnap(1)` manual page.

See [“About FileSnaps”](#) on page 580.

Online backup

VxFS provides online data backup using the snapshot feature. An image of a mounted file system instantly becomes an exact read-only copy of the file system at a specific point in time. The original file system is called the snapped file system, the copy is called the snapshot.

When changes are made to the snapped file system, the old data is copied to the snapshot. When the snapshot is read, data that has not changed is read from the snapped file system, changed data is read from the snapshot.

Backups require one of the following methods:

- Copying selected files from the snapshot file system (using `find` and `cpio`)
- Backing up the entire file system (using `fscat`)
- Initiating a full or incremental backup (using `vxdump`)

See [“About snapshot file systems”](#) on page 583.

Quotas

VxFS supports quotas, which allocate per-user and per-group quotas and limit the use of two principal resources: files and data blocks. You can assign quotas for each of these resources. Each quota consists of two limits for each resource: hard limit and soft limit.

The hard limit represents an absolute limit on data blocks or files. A user can never exceed the hard limit under any circumstances.

The soft limit is lower than the hard limit and can be exceeded for a limited amount of time. This allows users to exceed limits temporarily as long as they fall under those limits before the allotted time expires.

See [“About quota limits”](#) on page 923.

Support for databases

Databases are usually created on file systems to simplify backup, copying, and moving tasks and are slower compared to databases on raw disks.

Using Quick I/O for Databases feature with VxFS lets systems retain the benefits of having a database on a file system without sacrificing performance. Quick I/O creates regular, preallocated files to use as character devices. Databases can be created on the character devices to achieve the same performance as databases created on raw disks.

Treating regular VxFS files as raw devices has the following advantages for databases:

- Commercial database servers such as Oracle Server can issue kernel supported asynchronous I/O calls on these pseudo devices but not on regular files. Server can issue kernel supported asynchronous I/O calls on these pseudo devices but not on regular files.
- `read()` and `write()` system calls issued by the database server can avoid the acquisition and release of read/write locks inside the kernel that take place on regular files.

- VxFS can avoid double buffering of data already buffered by the database server. This ability frees up resources for other purposes and results in better performance.
- Since I/O to these devices bypasses the system buffer cache, VxFS saves on the cost of copying data between user space and kernel space when data is read from or written to a regular file. This process significantly reduces CPU time per I/O transaction compared to that of buffered I/O.

Cluster file systems

Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) allows clustered servers to mount and use a file system simultaneously as if all applications using the file system were running on the same server. The Veritas Volume Manager cluster functionality (CVM) makes logical volumes and raw device applications accessible through a cluster.

SFCFSHA uses a symmetric architecture in which all nodes in the cluster can simultaneously function as metadata servers. SFCFSHA still has some remnants of the old master/slave or primary/secondary concept. The first server to mount each cluster file system becomes its primary; all other nodes in the cluster become secondaries. Applications access the user data in files directly from the server on which they are running. Each SFCFSHA node has its own intent log. File system operations, such as allocating or deleting files, can originate from any node in the cluster.

Installing VxFS and enabling the cluster feature does not create a cluster file system configuration. File system clustering requires other Veritas products to enable communication services and provide storage resources. These products are packaged with VxFS in the Storage Foundation Cluster File System to provide a complete clustering environment.

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

To be a cluster mount, a file system must be mounted using the `mount -o cluster` option. File systems mounted without the `-o cluster` option are termed local mounts.

SFCFSHA functionality is separately licensed.

Cross-platform data sharing

Cross-platform data sharing (CDS) allows data to be serially shared among heterogeneous systems where each system has direct access to the physical

devices that hold the data. This feature can be used only in conjunction with Veritas Volume Manager (VxVM).

See the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

File Change Log

The VxFS File Change Log (FCL) tracks changes to files and directories in a file system. The File Change Log can be used by applications such as backup products, webcrawlers, search and indexing engines, and replication software that typically scan an entire file system searching for modifications since a previous scan. FCL functionality is a separately licensed feature.

See [“About File Change Log”](#) on page 930.

Reverse path name lookup

The reverse path name lookup feature obtains the full path name of a file or directory from the inode number of that file or directory. The reverse path name lookup feature can be useful for a variety of applications, such as for clients of the VxFS File Change Log feature, in backup and restore utilities, and for replication products. Typically, these applications store information by inode numbers because a path name for a file or directory can be very long, thus the need for an easy method of obtaining a path name.

See [“Reverse path name lookup”](#) on page 937.

Multi-volume file system support

The multi-volume support (MVS) feature allows several volumes to be represented by a single logical object. All I/O to and from an underlying logical volume is directed by way of volume sets. You can create a single VxFS file system on this multi-volume set. This feature can be used only in conjunction with VxVM. MVS functionality is a separately licensed feature.

See [“About multi-volume file systems”](#) on page 709.

SmartTier

The SmartTier option is built on multi-volume support technology. Using SmartTier, you can map more than one volume to a single file system. You can then configure policies that automatically relocate files from one volume to another, or relocate files by running file relocation commands. Having multiple volumes lets you determine where files are located, which can improve performance for applications

that access specific types of files. SmartTier functionality is a separately licensed feature.

Note: In the previous VxFS 5.x releases, SmartTier was known as Dynamic Storage Tiering.

See [“About SmartTier”](#) on page 724.

Thin Reclamation of a file system

Storage is allocated from a Thin Storage LUN when files are created and written to a file system. This storage is not given back to the Thin Storage LUN when a file is deleted or the file size is shrunk. As such, the file system must perform the explicit task of releasing the free storage to the Thin Storage LUN. This is performed by the Storage Foundation Thin Reclamation feature. Thin Reclamation is only supported on VxFS file systems mounted on a VxVM volume.

See [“Thin Reclamation”](#) on page 233.

Partitioned directories

Normally, a large volume of parallel threads performing access and updates on a directory that commonly exist in an file system suffers from exponentially longer wait times for the threads. This feature creates partitioned directories to improve the directory performance of file systems. When any directory crosses the tunable threshold, this feature takes an exclusive lock on the directory inode and redistributes the entries into various respective hash directories. These hash directories are not visible in the name-space view of the user or operating system. For every new create, delete, or lookup thread, this feature performs a lookup for the respective hashed directory (depending on the target name) and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance.

This feature operates only on disk layout Version 8 or later file systems.

See the `vxtunefs(1M)` and `fsadm_vxfs(1M)` manual pages.

Data deduplication on Solaris SPARC

On Solaris SPARC, you can perform post-process periodic deduplication in a file system to eliminate duplicate data without any continuous cost. You can verify whether data is duplicated on demand, and then efficiently and securely eliminate the duplicates. This feature requires an Enterprise license.

See [“About deduplicating data”](#) on page 788.

File compression

Compressing files reduces the space used by files, while retaining the accessibility of the files and being transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files. Reads cause data to be uncompressed in memory, only; the on-disk copy of the file remains compressed. In contrast, after a write, the new data is uncompressed on disk.

See [“About compressing files”](#) on page 799.

Veritas File System performance enhancements

Traditional file systems employ block-based allocation schemes that provide adequate random access and latency for small files, but which limit throughput for larger files. As a result, they are less than optimal for commercial environments.

Veritas File System (VxFS) addresses this file system performance issue through an alternative allocation method and increased user control over allocation, I/O, and caching policies.

See [“Using Veritas File System”](#) on page 89.

VxFS provides the following performance enhancements:

- Data synchronous I/O
- Direct I/O and discovered direct I/O
- [Delayed allocation for extending writes](#)
- [Enhanced I/O performance](#)
- Caching advisories
- Enhanced directory features
- Explicit file alignment, extent size, and preallocation controls
- Tunable I/O parameters
- Integration with Veritas Volume Manager (VxVM)
- Support for large directories

Note: VxFS reduces the file lookup time in directories with an extremely large number of files.

- Partitioned directories

Enhanced I/O performance

VxFS provides enhanced I/O performance by applying an aggressive I/O clustering policy, integrating with VxVM, and allowing application specific parameters to be set on a per-file system basis.

See [“Enhanced I/O clustering”](#) on page 88.

See [“Veritas Volume Manager integration with Veritas File System for enhanced I/O performance”](#) on page 88.

See [“Application-specific parameters for enhanced I/O performance”](#) on page 88.

Enhanced I/O clustering

I/O clustering is a technique of grouping multiple I/O operations together for improved performance. VxFS I/O policies provide more aggressive clustering processes than other file systems and offer higher I/O throughput when using large files. The resulting performance is comparable to that provided by raw disk.

Veritas Volume Manager integration with Veritas File System for enhanced I/O performance

VxFS interfaces with VxVM to determine the I/O characteristics of the underlying volume and perform I/O accordingly. VxFS also uses this information when using `mkfs` to perform proper allocation unit alignments for efficient I/O operations from the kernel.

As part of VxFS/VxVM integration, VxVM exports a set of I/O parameters to achieve better I/O performance. This interface can enhance performance for different volume configurations such as RAID-5, striped, and mirrored volumes. Full stripe writes are important in a RAID-5 volume for strong I/O performance. VxFS uses these parameters to issue appropriate I/O requests to VxVM.

Application-specific parameters for enhanced I/O performance

You can set application specific parameters on a per-file system basis to improve I/O performance.

- Discovered Direct I/O
All sizes above this value would be performed as direct I/O.
- Maximum Direct I/O Size
This value defines the maximum size of a single direct I/O.

See the `vxtunefs(1M)` and `tunefstab(4)` manual pages.

Delayed allocation for extending writes

Delayed allocation skips the allocations for extending writes and completes the allocations in a background thread. With this approach, Veritas File System (VxFS) performs a smaller number of large allocations instead of performing a large number of small allocations, which reduces the file system's fragmentation. Fast-moving temporary files do not have blocks allocated and thus do not add to the file system's fragmentation.

When a file is appended, the allocation to the file is skipped and the file is added to the delayed allocation list. The range for which the allocation is skipped is recorded in the inode. The `write()` system call returns immediately after the user pages are copied to the page cache. The actual allocations to the file occur when the scheduler thread picks the file for allocation. If the file is truncated or removed, allocations are not required.

Delayed allocation is turned on by default for extending writes. Delayed allocation is not dependent on the file system disk layout version. This feature does not require any `mount` options. You can turn off and turn on this feature by using the `vxtune fs` command. You can display the delayed allocation range in the file by using the `fsmmap` command.

See the `vxtune fs(1M)` and `fsmmap(1M)` manual pages.

For instances where the file data must be written to the disk immediately, delayed allocation is disabled on the file. The following are the examples of such instances: direct I/O, concurrent I/O, FDD/ODM access, and synchronous I/O. Delayed allocation is not supported on memory-mapped files, BSD quotas, and shared mount points in a Cluster File System (CFS). When BSD quotas are enabled on a file system, delayed allocation is turned off automatically for that file system.

Using Veritas File System

The following list contains the main methods to use, manage, modify, and tune VxFS:

- [Online system administration](#)
- [Application program interface](#)

Online system administration

Veritas File System (VxFS) provides command line interface (CLI) operations that are described throughout this guide and in manual pages.

VxFS allows you to run a number of administration tasks while the file system is online. Two of the more important tasks include:

- Defragmentation
- File system resizing

Defragmentation

Free resources are initially aligned and allocated to files in an order that provides optimal performance. On an active file system, the original order of free resources is lost over time as files are created, removed, and resized. The file system is spread farther along the disk, leaving unused gaps or fragments between areas that are in use. This process is known as fragmentation and leads to degraded performance because the file system has fewer options when assigning a free extent to a file (a group of contiguous data blocks).

VxFS provides the online administration utility `fsadm` to resolve the problem of fragmentation.

The `fsadm` utility defragments a mounted file system by performing the following actions:

- Removing unused space from directories
- Making all small files contiguous
- Consolidating free blocks for file system use

This utility can run on demand and should be scheduled regularly as a cron job.

File system resizing

A file system is assigned a specific size as soon as it is created; the file system may become too small or too large as changes in file system usage take place over time.

VxFS is capable of increasing or decreasing the file system size while in use. Many competing file systems can not do this. The VxFS utility `fsadm` can expand or shrink a file system without unmounting the file system or interrupting user productivity. However, to expand a file system, the underlying device on which it is mounted must be expandable.

VxVM facilitates expansion using virtual disks that can be increased in size while in use. The VxFS and VxVM packages complement each other to provide online expansion capability. Use the `vxresize` command when resizing both the volume and the file system. The `vxresize` command guarantees that the file system shrinks or grows along with the volume. You can also use the `vxassist` command combined with the `fsadm_vxfs` command for this purpose; however, Symantec recommends that you use the `vxresize` command instead.

See the `vxresize(1M)` manual page.

See [“Growing the existing storage by adding a new LUN”](#) on page 186.

Application program interface

Veritas File System Developer's Kit (SDK) provides developers with the information necessary to use the application programming interfaces (APIs) to modify and tune various features and components of File System.

See the *Veritas File System Programmer's Reference Guide*.

VxFS conforms to the System V Interface Definition (SVID) requirements and supports user access through the Network File System (NFS). Applications that require performance features not available with other file systems can take advantage of VxFS enhancements.

Expanded application facilities

VxFS provides API functions frequently associated with commercial applications that make it possible to perform the following actions:

- Preallocate space for a file
- Specify a fixed extent size for a file
- Bypass the system buffer cache for file I/O
- Specify the expected access pattern for a file

Because these functions are provided using VxFS-specific IOCTL system calls, most existing UNIX system applications do not use them. The VxFS-specific `cp`, `cpio`, and `mv` utilities use the functions to preserve extent attributes and allocate space more efficiently. The current attributes of a file can be listed using the `getext` or VxFS-specific `ls` command. The functions can also improve performance for custom applications. For portability reasons, these applications must check which file system type they are using before using these functions.

How Veritas Volume Manager works

This chapter includes the following topics:

- [How VxVM works with the operating system](#)
- [How VxVM handles storage management](#)
- [Volume layouts in VxVM](#)
- [Online relayout](#)
- [Volume resynchronization](#)
- [Hot-relocation](#)
- [Dirty region logging](#)
- [Volume snapshots](#)
- [FastResync](#)
- [Volume sets](#)
- [Configuration of volumes on SAN storage](#)

How VxVM works with the operating system

VxVM operates as a subsystem between your operating system and your data management systems, such as file systems and database management systems. VxVM is tightly coupled with the operating system. Before a disk or LUN can be brought under VxVM control, the disk must be accessible through the operating system device interface. VxVM is layered on top of the operating system interface services, and is dependent upon how the operating system accesses physical disks.

VxVM is dependent upon the operating system for the following functionality:

- operating system (disk) devices
- device handles
- VxVM Dynamic Multi-Pathing (DMP) metadvice

VxVM relies on the following constantly-running daemons and kernel threads for its operation:

<code>vxconfigd</code>	The VxVM configuration daemon maintains disk and group configurations and communicates configuration changes to the kernel, and modifies configuration information stored on disks.
<code>vxiod</code>	VxVM I/O kernel threads provide extended I/O operations without blocking calling processes. By default, 16 I/O threads are started at boot time, and at least one I/O thread must continue to run at all times.
<code>vxrelocd</code>	The hot-relocation daemon monitors VxVM for events that affect redundancy, and performs hot-relocation to restore redundancy. If thin provision disks are configured in the system, then the storage space of a deleted volume is reclaimed by this daemon as configured by the policy.

How data is stored

Several methods are used to store data on physical disks. These methods organize data on the disk so the data can be stored and retrieved efficiently. The basic method of disk organization is called formatting. Formatting prepares the hard disk so that files can be written to and retrieved from the disk by using a prearranged storage pattern.

Two methods are used to store information on formatted hard disks: physical-storage layout and logical-storage layout. VxVM uses the logical-storage layout method.

See [“How VxVM handles storage management”](#) on page 93.

How VxVM handles storage management

VxVM uses the following types of objects to handle storage management:

Physical objects	Physical disks, LUNs (virtual disks implemented in hardware), or other hardware with block and raw operating system device interfaces that are used to store data. See “Physical objects” on page 94.
Virtual objects	When one or more physical disks are brought under the control of VxVM, it creates virtual objects called volumes on those physical disks. Each volume records and retrieves data from one or more physical disks. Volumes are accessed by file systems, databases, or other applications in the same way that physical disks are accessed. Volumes are also composed of other virtual objects (plexes and subdisks) that are used in changing the volume configuration. Volumes and their virtual components are called virtual objects or VxVM objects. See “Virtual objects” on page 96.

Physical objects

A physical disk is the basic storage device (media) where the data is ultimately stored. You can access the data on a physical disk by using a device name to locate the disk. The physical disk device name varies with the computer system you use. Not all parameters are used on all systems.

Typical device names are of the form `c#t#d#s#`, where `c#` specifies the controller, `t#` specifies the target ID, `d#` specifies the disk, and `s#` specifies the partition or slice. For example, device name `c0t0d0s2` is the entire hard disk connected to controller number 0 in the system, with a target ID of 0, and physical disk number 0.

[Figure 4-1](#) shows how a physical disk and device name (*devname*) are illustrated in this document.

Figure 4-1 Physical disk example




VxVM writes identification information on physical disks under VxVM control (VM disks). VxVM disks can be identified even after physical disk disconnection or system outages. VxVM can then re-form disk groups and logical objects to provide failure detection and to speed system recovery.

Partitions

Figure 4-2 shows how a physical disk can be divided into one or more slices, also known as partitions.

Figure 4-2 Slice example



The slice number is added at the end of the *devname*, and is denoted by `s#`. Note that slice `s2` refers to an entire physical disk for disks that are not Extensible Firmware Interface (EFI) disks.


Disk arrays

Performing I/O to disks is a relatively slow process because disks are physical devices that require time to move the heads to the correct position on the disk before reading or writing. If all of the read or write operations are done to individual disks, one at a time, the read-write time can become unmanageable. Performing these operations on multiple disks can help to reduce this problem.

A disk array is a collection of physical disks that VxVM can represent to the operating system as one or more virtual disks or volumes. The volumes created by VxVM look and act to the operating system like physical disks. Applications that interact with volumes should work in the same way as with physical disks.

Figure 4-3 shows how VxVM represents the disks in a disk array as several volumes to the operating system.

Figure 4-3 How VxVM presents the disks in a disk array as volumes to the operating system



Data can be spread across several disks within an array, or across disks spanning multiple arrays, to distribute or balance I/O operations across the disks. Using parallel I/O across multiple disks in this way improves I/O performance by increasing data transfer speed and overall throughput for the array.

Virtual objects

VxVM uses multiple virtualization layers to provide distinct functionality and reduce physical limitations.

Virtual objects in VxVM include the following:

- Disk groups
See “[Disk groups](#)” on page 98.
- VM disks
See “[VM disks](#)” on page 99.
- Subdisks
See “[Subdisks](#)” on page 100.
- Plexes
See “[Plexes](#)” on page 101.
- Volumes
See “[Volumes](#)” on page 102.

The connection between physical objects and VxVM objects is made when you place a physical disk under VxVM control.

After installing VxVM on a host system, you must bring the contents of physical disks under VxVM control by collecting the VM disks into disk groups and allocating the disk group space to create logical volumes.

Bringing the contents of physical disks under VxVM control is accomplished only if VxVM takes control of the physical disks and the disk is not under control of another storage manager such as Sun Microsystems Solaris Volume Manager software.

For more information on how Solaris Volume Manager and VM disks co-exist or how to convert LVM disks to VM disks, see the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

VxVM creates virtual objects and makes logical connections between the objects. The virtual objects are then used by VxVM to do storage management tasks.

The `vxprint` command displays detailed information about the VxVM objects that exist on a system.


See the `vxprint(1M)` manual page.

Combining virtual objects in VxVM

VxVM virtual objects are combined to build volumes. The virtual objects contained in volumes are VM disks, disk groups, subdisks, and plexes. VxVM virtual objects are organized in the following ways:

- VM disks are grouped into disk groups
- Subdisks (each representing a specific region of a disk) are combined to form plexes
- Volumes are composed of one or more plexes

[Figure 4-4](#) shows the connections between Veritas Volume Manager virtual objects and how they relate to physical disks.

Figure 4-4 Connection between objects in VxVM

The disk group contains three VM disks which are used to create two volumes. Volume `vol01` is simple and has a single plex. Volume `vol02` is a mirrored volume with two plexes.

The various types of virtual objects (disk groups, VM disks, subdisks, plexes and volumes) are described in the following sections. Other types of objects exist in Veritas Volume Manager, such as data change objects (DCOs), and volume sets, to provide extended functionality.

Disk groups

A disk group is a collection of disks that share a common configuration and which are managed by VxVM. A disk group configuration is a set of records with detailed information about related VxVM objects, their attributes, and their connections. A disk group name can be up to 31 characters long.

See [“VM disks”](#) on page 99.

In releases before VxVM 4.0, the default disk group was `rootdg` (the root disk group). For VxVM to function, the `rootdg` disk group had to exist and it had to contain at least one disk. This requirement no longer exists. VxVM can work without any disk groups configured (although you must set up at least one disk group before you can create any volumes of other VxVM objects).

You can create additional disk groups when you need them. Disk groups allow you to group disks into logical collections. A disk group and its components can be moved as a unit from one host machine to another.

See [“Reorganizing the contents of disk groups”](#) on page 814.

Volumes are created within a disk group. A given volume and its plexes and subdisks must be configured from disks in the same disk group.

VM disks


When you place a physical disk under VxVM control, a VM disk is assigned to the physical disk. A VM disk is under VxVM control and is usually in a disk group. Each VM disk corresponds to at least one physical disk or disk partition. VxVM allocates storage from a contiguous area of VxVM disk space.

A VM disk typically includes a public region (allocated storage) and a small private region where VxVM internal configuration information is stored.

Each VM disk has a unique disk media name (a virtual disk name). You can either define a disk name of up to 31 characters, or allow VxVM to assign a default name that takes the form `diskgroup##`, where *diskgroup* is the name of the disk group to which the disk belongs.

See [“Disk groups”](#) on page 98.

[Figure 4-5](#) shows a VM disk with a media name of `disk01` that is assigned to the physical disk, *devname*.

Figure 4-5 VM disk example


Subdisks

A subdisk is a set of contiguous disk blocks. A block is a unit of space on the disk. VxVM allocates disk space using subdisks. A VM disk can be divided into one or more subdisks. Each subdisk represents a specific portion of a VM disk, which is mapped to a specific region of a physical disk.

The default name for a VM disk is `diskgroup##` and the default name for a subdisk is `diskgroup##-##`, where *diskgroup* is the name of the disk group to which the disk belongs.

See “Disk groups” on page 98.


Figure 4-6 shows `disk01-01` is the name of the first subdisk on the VM disk named `disk01`.

Figure 4-6 Subdisk example

A VM disk can contain multiple subdisks, but subdisks cannot overlap or share the same portions of a VM disk. To ensure integrity, VxVM rejects any commands that try to create overlapping subdisks.

Figure 4-7 shows a VM disk with three subdisks, which are assigned from one physical disk.

Figure 4-7 Example of three subdisks assigned to one VM Disk




Any VM disk space that is not part of a subdisk is free space. You can use free space to create new subdisks.

Plexes

VxVM uses subdisks to build virtual objects called plexes. A plex consists of one or more subdisks located on one or more physical disks.

Figure 4-8 shows an example of a plex with two subdisks.

Figure 4-8 Example of a plex with two subdisks



You can organize data on subdisks to form a plex by using the following methods:

- concatenation
- striping (RAID-0)
- mirroring (RAID-1)
- striping with parity (RAID-5)

Concatenation, striping (RAID-0), mirroring (RAID-1) and RAID-5 are types of volume layout.

See [“Volume layouts in VxVM”](#) on page 104.

Volumes

A volume is a virtual disk device that appears to applications, databases, and file systems like a physical disk device, but does not have the physical limitations of a physical disk device. A volume consists of one or more plexes, each holding a copy of the selected data in the volume. Due to its virtual nature, a volume is not restricted to a particular disk or a specific area of a disk. The configuration of a volume can be changed by using VxVM user interfaces. Configuration changes can be accomplished without causing disruption to applications or file systems that are using the volume. For example, a volume can be mirrored on separate disks or moved to use different disk storage.


VxVM uses the default naming conventions of `vol##` for volumes and `vol##-##` for plexes in a volume. For ease of administration, you can choose to select more meaningful names for the volumes that you create.

A volume may be created under the following constraints:

- Its name can contain up to 31 characters.
- It can consist of up to 32 plexes, each of which contains one or more subdisks.
- It must have at least one associated plex that has a complete copy of the data in the volume with at least one associated subdisk.
- All subdisks within a volume must belong to the same disk group.

Figure 4-9 shows a volume `vol01` with a single plex.


Figure 4-9 Example of a volume with one plex



The volume `vol01` has the following characteristics:

- It contains one plex named `vol01-01`.
- The plex contains one subdisk named `disk01-01`.
- The subdisk `disk01-01` is allocated from VM disk `disk01`.

Figure 4-10 shows a mirrored volume `vol06` with two data plexes.

Figure 4-10 Example of a volume with two plexes

Each plex of the mirror contains a complete copy of the volume data.

The volume `vol06` has the following characteristics:

- It contains two plexes named `vol06-01` and `vol06-02`.
- Each plex contains one subdisk.
- Each subdisk is allocated from a different VM disk (`disk01` and `disk02`).

See [“Mirroring \(RAID-1\)”](#) on page 111.

VxVM supports the concept of layered volumes in which subdisks can contain volumes.

See [“Layered volumes”](#) on page 119.

About the configuration daemon in VxVM

The VxVM configuration daemon (`vxconfigd`) provides the interface between VxVM commands and the kernel device drivers. `vxconfigd` handles configuration change requests from VxVM utilities, communicates the change requests to the VxVM kernel, and modifies configuration information stored on disk. `vxconfigd` also initializes VxVM when the system is booted.

The `vxctl` command is the command-line interface to the `vxconfigd` daemon.

You can use `vxctl` to:

- Control the operation of the `vxconfigd` daemon.
- Change the system-wide definition of the default disk group.

In VxVM 4.0 and later releases, disk access records are no longer stored in the `/etc/vx/volboot` file. Non-persistent disk access records are created by scanning the disks at system startup. Persistent disk access records for `simple` and `nopriv` disks are permanently stored in the `/etc/vx/darecs` file in the `root` file system.

The `vxconfigd` daemon reads the contents of this file to locate the disks and the configuration databases for their disk groups.

The `/etc/vx/darecs` file is also used to store definitions of foreign devices that are not autoconfigurable. Such entries may be added by using the `vxddladm addforeign` command.

See the `vxddladm(1M)` manual page.

If your system is configured to use Dynamic Multi-Pathing (DMP), you can also use `vxctl` to:

- Reconfigure the DMP database to include disk devices newly attached to, or removed from the system.
- Create DMP device nodes in the `/dev/vx/dmp` and `/dev/vx/rdmp` directories.
- Update the DMP database with changes in path type for active/passive disk arrays. Use the utilities provided by the disk-array vendor to change the path type between primary and secondary.

See the `vxctl(1M)` manual page.

Multiple paths to disk arrays

Some disk arrays provide multiple ports to access their disk devices. These ports, coupled with the host bus adaptor (HBA) controller and any data bus or I/O processor local to the array, make up multiple hardware paths to access the disk devices. Such disk arrays are called multipathed disk arrays. This type of disk array can be connected to host systems in many different configurations, (such as multiple ports connected to different controllers on a single host, chaining of the ports through a single controller on a host, or ports connected to different hosts simultaneously).

See [“How DMP works”](#) on page 170.

Volume layouts in VxVM

A VxVM virtual device is defined by a volume. A volume has a layout defined by the association of a volume to one or more plexes, each of which map to one or more subdisks. The volume presents a virtual device interface that is exposed to other applications for data access. These logical building blocks re-map the volume address space through which I/O is re-directed at run-time.

Different volume layouts provide different levels of availability and performance. A volume layout can be configured and changed to provide the desired level of service.

Non-layered volumes

In a non-layered volume, a subdisk maps directly to a VM disk. This allows the subdisk to define a contiguous extent of storage space backed by the public region of a VM disk. When active, the VM disk is directly associated with an underlying physical disk. The combination of a volume layout and the physical disks therefore determines the storage service available from a given virtual device.

Layered volumes

A layered volume is constructed by mapping its subdisks to underlying volumes. The subdisks in the underlying volumes must map to VM disks, and hence to attached physical storage.

Layered volumes allow for more combinations of logical compositions, some of which may be desirable for configuring a virtual device. For example, layered volumes allow for high availability when using striping. Because permitting free use of layered volumes throughout the command level would have resulted in unwieldy administration, some ready-made layered volume configurations are designed into VxVM.

See [“Layered volumes”](#) on page 119.

These ready-made configurations operate with built-in rules to automatically match desired levels of service within specified constraints. The automatic configuration is done on a “best-effort” basis for the current command invocation working against the current configuration.

To achieve the desired storage service from a set of virtual devices, it may be necessary to include an appropriate set of VM disks into a disk group and to execute multiple configuration commands.

To the extent that it can, VxVM handles initial configuration and on-line re-configuration with its set of layouts and administration interface to make this job easier and more deterministic.

Layout methods

Data in virtual objects is organized to create volumes by using the following layout methods:

- Concatenation, spanning, and carving
See [“Concatenation, spanning, and carving”](#) on page 106.
- Striping (RAID-0)
See [“Striping \(RAID-0\)”](#) on page 108.
- Mirroring (RAID-1)

- See [“Mirroring \(RAID-1\)”](#) on page 111.
- Striping plus mirroring (mirrored-stripe or RAID-0+1)
See [“Striping plus mirroring \(mirrored-stripe or RAID-0+1\)”](#) on page 112.
- Mirroring plus striping (striped-mirror, RAID-1+0 or RAID-10)
See [“Mirroring plus striping \(striped-mirror, RAID-1+0 or RAID-10\)”](#) on page 113.
- RAID-5 (striping with parity)
See [“RAID-5 \(striping with parity\)”](#) on page 114.


Concatenation, spanning, and carving

Concatenation maps data in a linear manner onto one or more subdisks in a plex. To access all of the data in a concatenated plex sequentially, data is first accessed in the first subdisk from the beginning to the end. Data is then accessed in the remaining subdisks sequentially from the beginning to the end of each subdisk, until the end of the last subdisk.

The subdisks in a concatenated plex do not have to be physically contiguous and can belong to more than one VM disk. Concatenation using subdisks that reside on more than one VM disk is called spanning.

[Figure 4-11](#) shows the concatenation of two subdisks from the same VM disk.

If a single LUN or disk is split into multiple subdisks, and each subdisk belongs to a unique volume, it is called carving.


Figure 4-11 Example of concatenation

The blocks n , $n+1$, $n+2$ and $n+3$ (numbered relative to the start of the plex) are contiguous on the plex, but actually come from two distinct subdisks on the same physical disk.

The remaining free space in the subdisk `disk01-02` on VM disk `disk01` can be put to other uses.

You can use concatenation with multiple subdisks when there is insufficient contiguous space for the plex on any one disk. This form of concatenation can be used for load balancing between disks, and for head movement optimization on a particular disk.

Figure 4-12 shows data spread over two subdisks in a spanned plex.

Figure 4-12 Example of spanning

The blocks `n`, `n+1`, `n+2` and `n+3` (numbered relative to the start of the plex) are contiguous on the plex, but actually come from two distinct subdisks from two distinct physical disks.

The remaining free space in the subdisk `disk02-02` on VM disk `disk02` can be put to other uses.

Warning: Spanning a plex across multiple disks increases the chance that a disk failure results in failure of the assigned volume. Use mirroring or RAID-5 to reduce the risk that a single disk failure results in a volume failure.

Striping (RAID-0)

Striping (RAID-0) is useful if you need large amounts of data written to or read from physical disks, and performance is important. Striping is also helpful in balancing the I/O load from multi-user applications across multiple disks. By using parallel data transfer to and from multiple disks, striping significantly improves data-access performance.

Striping maps data so that the data is interleaved among two or more physical disks. A striped plex contains two or more subdisks, spread out over two or more physical disks. Data is allocated alternately and evenly to the subdisks of a striped plex.


The subdisks are grouped into “columns,” with each physical disk limited to one column. Each column contains one or more subdisks and can be derived from one or more physical disks. The number and sizes of subdisks per column can vary. Additional subdisks can be added to columns, as necessary.

Warning: Striping a volume, or splitting a volume across multiple disks, increases the chance that a disk failure will result in failure of that volume.

If five volumes are striped across the same five disks, then failure of any one of the five disks will require that all five volumes be restored from a backup. If each volume is on a separate disk, only one volume has to be restored. (As an alternative to or in conjunction with striping, use mirroring or RAID-5 to substantially reduce the chance that a single disk failure results in failure of a large number of volumes.)

Data is allocated in equal-sized stripe units that are interleaved between the columns. Each stripe unit is a set of contiguous blocks on a disk. The default stripe unit size is 64 kilobytes.

Figure 4-13 shows an example with three columns in a striped plex, six stripe units, and data striped over the three columns.



A stripe consists of the set of stripe units at the same positions across all columns. In the figure, stripe units 1, 2, and 3 constitute a single stripe.

Viewed in sequence, the first stripe consists of:

- stripe unit 1 in column 0
- stripe unit 2 in column 1
- stripe unit 3 in column 2


The second stripe consists of:

- stripe unit 4 in column 0
- stripe unit 5 in column 1
- stripe unit 6 in column 2

Striping continues for the length of the columns (if all columns are the same length), or until the end of the shortest column is reached. Any space remaining at the end of subdisks in longer columns becomes unused space.

Figure 4-14 shows a striped plex with three equal sized, single-subdisk columns.

Figure 4-14 Example of a striped plex with one subdisk per column




There is one column per physical disk. This example shows three subdisks that occupy all of the space on the VM disks. It is also possible for each subdisk in a

striped plex to occupy only a portion of the VM disk, which leaves free space for other disk management tasks.

Figure 4-15 shows a striped plex with three columns containing subdisks of different sizes.

Figure 4-15 Example of a striped plex with concatenated subdisks per column



Each column contains a different number of subdisks. There is one column per physical disk. Striped plexes can be created by using a single subdisk from each of the VM disks being striped across. It is also possible to allocate space from different regions of the same disk or from another disk (for example, if the size of the plex is increased). Columns can also contain subdisks from different VM disks.

See [“Creating a striped volume”](#) on page 201.

Mirroring (RAID-1)

Mirroring uses multiple mirrors (plexes) to duplicate the information contained in a volume. In the event of a physical disk failure, the plex on the failed disk becomes

unavailable, but the system continues to operate using the unaffected mirrors. Similarly, mirroring two LUNs from two separate controllers lets the system operate if there is a controller failure.

Although a volume can have a single plex, at least two plexes are required to provide redundancy of data. Each of these plexes must contain disk space from different disks to achieve redundancy.

When striping or spanning across a large number of disks, failure of any one of those disks can make the entire plex unusable. Because the likelihood of one out of several disks failing is reasonably high, you should consider mirroring to improve the reliability (and availability) of a striped or spanned volume.

See [“Creating a mirrored volume”](#) on page 199.


Striping plus mirroring (mirrored-stripe or RAID-0+1)

VxVM supports the combination of mirroring above striping. The combined layout is called a mirrored-stripe layout. A mirrored-stripe layout offers the dual benefits of striping to spread data across multiple disks, while mirroring provides redundancy of data.

For mirroring above striping to be effective, the striped plex and its mirrors must be allocated from separate disks.

[Figure 4-16](#) shows an example where two plexes, each striped across three disks, are attached as mirrors to the same volume to create a mirrored-stripe volume.

Figure 4-16 Mirrored-stripe volume laid out on six disks



See [“Creating a mirrored-stripe volume”](#) on page 202.

The layout type of the data plexes in a mirror can be concatenated or striped. Even if only one is striped, the volume is still termed a mirrored-stripe volume. If they are all concatenated, the volume is termed a mirrored-concatenated volume.

Mirroring plus striping (striped-mirror, RAID-1+0 or RAID-10)

VxVM supports the combination of striping above mirroring. This combined layout is called a striped-mirror layout. Putting mirroring below striping mirrors each column of the stripe. If there are multiple subdisks per column, each subdisk can be mirrored individually instead of each column.


A striped-mirror volume is an example of a layered volume.

See [“Layered volumes”](#) on page 119.

As for a mirrored-stripe volume, a striped-mirror volume offers the dual benefits of striping to spread data across multiple disks, while mirroring provides redundancy of data. In addition, it enhances redundancy, and reduces recovery time after disk failure.

[Figure 4-17](#) shows an example where a striped-mirror volume is created by using each of three existing 2-disk mirrored volumes to form a separate column within a striped plex.


Figure 4-17 Striped-mirror volume laid out on six disks



See [“Creating a striped-mirror volume”](#) on page 202.

[Figure 4-18](#) shows that the failure of a disk in a mirrored-stripe layout detaches an entire data plex, thereby losing redundancy on the entire volume.

Figure 4-18 How the failure of a single disk affects mirrored-stripe and striped-mirror volumes



When the disk is replaced, the entire plex must be brought up to date. Recovering the entire plex can take a substantial amount of time. If a disk fails in a striped-mirror layout, only the failing subdisk must be detached, and only that portion of the volume loses redundancy. When the disk is replaced, only a portion of the volume needs to be recovered. Additionally, a mirrored-stripe volume is more vulnerable to being put out of use altogether should a second disk fail before the first failed disk has been replaced, either manually or by hot-relocation.

Compared to mirrored-stripe volumes, striped-mirror volumes are more tolerant of disk failure, and recovery time is shorter.

If the layered volume concatenates instead of striping the underlying mirrored volumes, the volume is termed a concatenated-mirror volume.

RAID-5 (striping with parity)


Although both mirroring (RAID-1) and RAID-5 provide redundancy of data, they use different methods. Mirroring provides data redundancy by maintaining multiple complete copies of the data in a volume. Data being written to a mirrored volume

is reflected in all copies. If a portion of a mirrored volume fails, the system continues to use the other copies of the data.

RAID-5 provides data redundancy by using parity. Parity is a calculated value used to reconstruct data after a failure. While data is being written to a RAID-5 volume, parity is calculated by doing an exclusive OR (XOR) procedure on the data. The resulting parity is then written to the volume. The data and calculated parity are contained in a plex that is “striped” across multiple disks. If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be recreated from the remaining data and parity information. It is also possible to mix concatenation and striping in the layout.

Figure 4-19 shows parity locations in a RAID-5 array configuration.

Figure 4-19 Parity locations in a RAID-5 model



Every stripe has a column containing a parity stripe unit and columns containing data. The parity is spread over all of the disks in the array, reducing the write time for large independent writes because the writes do not have to wait until a single parity disk can accept the data.

RAID-5 volumes can additionally perform logging to minimize recovery time. RAID-5 volumes use RAID-5 logs to keep a copy of the data and parity currently being written. RAID-5 logging is optional and can be created along with RAID-5 volumes or added later.

See “[Veritas Volume Manager RAID-5 arrays](#)” on page 116.


Note: VxVM supports RAID-5 for private disk groups, but not for shareable disk groups in a CVM environment. In addition, VxVM does not support the mirroring of RAID-5 volumes that are configured using Veritas Volume Manager software. RAID-5 LUNs hardware may be mirrored.

Traditional RAID-5 arrays

A traditional RAID-5 array is several disks organized in rows and columns. A column is a number of disks located in the same ordinal position in the array. A row is the minimal number of disks necessary to support the full width of a parity stripe.

Figure 4-20 shows the row and column arrangement of a traditional RAID-5 array.

Figure 4-20 Traditional RAID-5 array




This traditional array structure supports growth by adding more rows per column. Striping is accomplished by applying the first stripe across the disks in Row 0, then the second stripe across the disks in Row 1, then the third stripe across the Row 0 disks, and so on. This type of array requires all disks columns and rows to be of equal size.

Veritas Volume Manager RAID-5 arrays

The RAID-5 array structure in Veritas Volume Manager differs from the traditional structure. Due to the virtual nature of its disks and other objects, VxVM does not use rows.

Figure 4-21 shows how VxVM uses columns consisting of variable length subdisks, where each subdisk represents a specific area of a disk.

Figure 4-21 Veritas Volume Manager RAID-5 array

VxVM allows each column of a RAID-5 plex to consist of a different number of subdisks. The subdisks in a given column can be derived from different physical disks. Additional subdisks can be added to the columns as necessary. Striping is implemented by applying the first stripe across each subdisk at the top of each column, then applying another stripe below that, and so on for the length of the columns. Equal-sized stripe units are used for each column. For RAID-5, the default stripe unit size is 16 kilobytes.

See [“Striping \(RAID-0\)”](#) on page 108.

Note: Mirroring of RAID-5 volumes is not supported.

See [“Creating a RAID-5 volume”](#) on page 203.

Left-symmetric layout


There are several layouts for data and parity that can be used in the setup of a RAID-5 array. The implementation of RAID-5 in VxVM uses a left-symmetric layout. This provides optimal performance for both random I/O operations and large sequential I/O operations. However, the layout selection is not as critical for performance as are the number of columns and the stripe unit size.

Left-symmetric layout stripes both data and parity across columns, placing the parity in a different column for every stripe of data. The first parity stripe unit is located in the rightmost column of the first stripe. Each successive parity stripe unit is located

in the next stripe, shifted left one column from the previous parity stripe unit location. If there are more stripes than columns, the parity stripe unit placement begins in the rightmost column again.

Figure 4-22 shows a left-symmetric parity layout with five disks (one per column).

Figure 4-22 Left-symmetric layout



For each stripe, data is organized starting to the right of the parity stripe unit. In the figure, data organization for the first stripe begins at P0 and continues to stripe units 0-3. Data organization for the second stripe begins at P1, then continues to stripe unit 4, and on to stripe units 5-7. Data organization proceeds in this manner for the remaining stripes.

Each parity stripe unit contains the result of an exclusive OR (XOR) operation performed on the data in the data stripe units within the same stripe. If one column's data is inaccessible due to hardware or software failure, the data for each stripe can be restored by XORing the contents of the remaining columns data stripe units against their respective parity stripe units.

For example, if a disk corresponding to the whole or part of the far left column fails, the volume is placed in a degraded mode. While in degraded mode, the data from the failed column can be recreated by XORing stripe units 1-3 against parity stripe unit P0 to recreate stripe unit 0, then XORing stripe units 4, 6, and 7 against parity stripe unit P1 to recreate stripe unit 5, and so on.

Failure of more than one column in a RAID-5 plex detaches the volume. The volume is no longer allowed to satisfy read or write requests. Once the failed columns have been recovered, it may be necessary to recover user data from backups.


RAID-5 logging

Logging is used to prevent corruption of data during recovery by immediately recording changes to data and parity to a log area on a persistent device such as a volume on disk or in non-volatile RAM. The new data and parity are then written to the disks.

Without logging, it is possible for data not involved in any active writes to be lost or silently corrupted if both a disk in a RAID-5 volume and the system fail. If this double-failure occurs, there is no way of knowing if the data being written to the data portions of the disks or the parity being written to the parity portions have actually been written. Therefore, the recovery of the corrupted disk may be corrupted itself.

Figure 4-23 shows a RAID-5 volume configured across three disks (A, B and C).

Figure 4-23 Incomplete write to a RAID-5 volume



In this volume, recovery of disk B's corrupted data depends on disk A's data and disk C's parity both being complete. However, only the data write to disk A is complete. The parity write to disk C is incomplete, which would cause the data on disk B to be reconstructed incorrectly.


This failure can be avoided by logging all data and parity writes before committing them to the array. In this way, the log can be replayed, causing the data and parity updates to be completed before the reconstruction of the failed drive takes place.

Logs are associated with a RAID-5 volume by being attached as log plexes. More than one log plex can exist for each RAID-5 volume, in which case the log areas are mirrored.

Layered volumes

A layered volume is a virtual Veritas Volume Manager object that is built on top of other volumes. The layered volume structure tolerates failure better and has greater redundancy than the standard volume structure. For example, in a striped-mirror layered volume, each mirror (plex) covers a smaller area of storage space, so recovery is quicker than with a standard mirrored volume.

Figure 4-24 shows a typical striped-mirror layered volume where each column is represented by a subdisk that is built from an underlying mirrored volume.



The volume and striped plex in the “Managed by User” area allow you to perform normal tasks in VxVM. User tasks can be performed only on the top-level volume of a layered volume.

Underlying volumes in the “Managed by VxVM” area are used exclusively by VxVM and are not designed for user manipulation. You cannot detach a layered volume or perform any other operation on the underlying volumes by manipulating the internal structure. You can perform all necessary operations in the “Managed by User” area that includes the top-level volume and striped plex (for example, resizing the volume, changing the column width, or adding a column).

System administrators can manipulate the layered volume structure for troubleshooting or other operations (for example, to place data on specific disks). Layered volumes are used by VxVM to perform the following tasks and operations:

Creating striped-mirrors	See “Creating a striped-mirror volume” on page 202. See the <code>vxassist(1M)</code> manual page.
Creating concatenated-mirrors	See “Creating a concatenated-mirror volume” on page 201. See the <code>vxassist(1M)</code> manual page.
Online Relayout	See “Online relayout” on page 121. See the <code>vxassist(1M)</code> manual page. See the <code>vxrelayout(1M)</code> manual page.
Moving RAID-5 subdisks	See the <code>vxsd(1M)</code> manual page.
Creating Snapshots	See “Volume snapshots” on page 130. See the <code>vxassist(1M)</code> manual page. See the <code>vxsnap(1M)</code> manual page.

Online relayout

Online relayout allows you to convert between storage layouts in VxVM, with uninterrupted data access. Typically, you would do this to change the redundancy or performance characteristics of a volume. VxVM adds redundancy to storage either by duplicating the data (mirroring) or by adding parity (RAID-5). Performance characteristics of storage in VxVM can be changed by changing the striping parameters, which are the number of columns and the stripe width.

See [“Performing online relayout”](#) on page 831.

How online relayout works

Online relayout allows you to change the storage layouts that you have already created in place without disturbing data access. You can change the performance characteristics of a particular layout to suit your changed requirements. You can transform one layout to another by invoking a single command.

For example, if a striped layout with a 128KB stripe unit size is not providing optimal performance, you can use relayout to change the stripe unit size.

File systems mounted on the volumes do not need to be unmounted to achieve this transformation provided that the file system (such as Veritas File System) supports online shrink and grow operations.

Online relayout reuses the existing storage space and has space allocation policies to address the needs of the new layout. The layout transformation process converts

a given volume to the destination layout by using minimal temporary space that is available in the disk group.

The transformation is done by moving one portion of data at a time in the source layout to the destination layout. Data is copied from the source volume to the temporary area, and data is removed from the source volume storage area in portions. The source volume storage area is then transformed to the new layout, and the data saved in the temporary area is written back to the new layout. This operation is repeated until all the storage and data in the source volume has been transformed to the new layout.

The default size of the temporary area used during the relayout depends on the size of the volume and the type of relayout. For volumes larger than 50MB, the amount of temporary space that is required is usually 10% of the size of the volume, from a minimum of 50MB up to a maximum of 1GB. For volumes smaller than 50MB, the temporary space required is the same as the size of the volume.

The following error message displays the number of blocks required if there is insufficient free space available in the disk group for the temporary area:

```
tmpsize too small to perform this relayout (nblks minimum required)
```


You can override the default size used for the temporary area by using the `tmpsize` attribute to `vxassist`.

See the `vxassist(1M)` manual page.

As well as the temporary area, space is required for a temporary intermediate volume when increasing the column length of a striped volume. The amount of space required is the difference between the column lengths of the target and source volumes. For example, 20GB of temporary additional space is required to relayout a 150GB striped volume with 5 columns of length 30GB as 3 columns of length 50GB. In some cases, the amount of temporary space that is required is relatively large. For example, a relayout of a 150GB striped volume with 5 columns as a concatenated volume (with effectively one column) requires 120GB of space for the intermediate volume.

Additional permanent disk space may be required for the destination volumes, depending on the type of relayout that you are performing. This may happen, for example, if you change the number of columns in a striped volume.

[Figure 4-25](#) shows how decreasing the number of columns can require disks to be added to a volume.


Figure 4-25 Example of decreasing the number of columns in a volume

Note that the size of the volume remains the same but an extra disk is needed to extend one of the columns.

The following are examples of operations that you can perform using online relayout:

- Remove parity from a RAID-5 volume to change it to a concatenated, striped, or layered volume.


[Figure 4-26](#) shows an example of applying relayout a RAID-5 volume.

Figure 4-26 Example of relayout of a RAID-5 volume to a striped volume

Note that removing parity decreases the overall storage space that the volume requires.

- Add parity to a volume to change it to a RAID-5 volume.

[Figure 4-27](#) shows an example.


Figure 4-27 Example of relayout of a concatenated volume to a RAID-5 volume

Note that adding parity increases the overall storage space that the volume requires.

- Change the number of columns in a volume.

Figure 4-28 shows an example of changing the number of columns.

Figure 4-28 Example of increasing the number of columns in a volume




Note that the length of the columns is reduced to conserve the size of the volume.

- Change the column stripe width in a volume.

Figure 4-29 shows an example of changing the column stripe width.

Figure 4-29 Example of increasing the stripe width for the columns in a volume



See [“Performing online relayout”](#) on page 831.

See [“Permitted relayout transformations”](#) on page 832.

Limitations of online relayout

Note the following limitations of online relayout:

- Log plexes cannot be transformed.
- Volume snapshots cannot be taken when there is an online relayout operation running on the volume.
- Online relayout cannot create a non-layered mirrored volume in a single step. It always creates a layered mirrored volume even if you specify a non-layered mirrored layout, such as `mirror-stripe` or `mirror-concat`. Use the `vxassist convert` command to turn the layered mirrored volume that results from a relayout into a non-layered volume.
- The usual restrictions apply for the minimum number of physical disks that are required to create the destination layout. For example, mirrored volumes require

at least as many disks as mirrors, striped and RAID-5 volumes require at least as many disks as columns, and striped-mirror volumes require at least as many disks as columns multiplied by mirrors.

- To be eligible for layout transformation, the plexes in a mirrored volume must have identical stripe widths and numbers of columns. Relayout is not possible unless you make the layouts of the individual plexes identical.
- Online relayout cannot transform sparse plexes, nor can it make any plex sparse. (A sparse plex is a plex that is not the same size as the volume, or that has regions that are not mapped to any subdisk.)
- The number of mirrors in a mirrored volume cannot be changed using relayout. Instead, use alternative commands, such as the `vxassist mirror` command.
- Only one relayout may be applied to a volume at a time.

Transformation characteristics

Transformation of data from one layout to another involves rearrangement of data in the existing layout to the new layout. During the transformation, online relayout retains data redundancy by mirroring any temporary space used. Read and write access to data is not interrupted during the transformation.

Data is not corrupted if the system fails during a transformation. The transformation continues after the system is restored and both read and write access are maintained.

You can reverse the layout transformation process at any time, but the data may not be returned to the exact previous storage location. Before you reverse a transformation that is in process, you must stop it.

You can determine the transformation direction by using the `vxrelayout status volume` command.

These transformations are protected against I/O failures if there is sufficient redundancy and space to move the data.

Transformations and volume length

Some layout transformations can cause the volume length to increase or decrease. If either of these conditions occurs, online relayout uses the `vxresize` command to shrink or grow a file system.

Volume resynchronization

When storing data redundantly and using mirrored or RAID-5 volumes, VxVM ensures that all copies of the data match exactly. However, under certain conditions (usually due to complete system failures), some redundant data on a volume can become inconsistent or unsynchronized. The mirrored data is not exactly the same as the original data. Except for normal configuration changes (such as detaching and reattaching a plex), this can only occur when a system crashes while data is being written to a volume.

Data is written to the mirrors of a volume in parallel, as is the data and parity in a RAID-5 volume. If a system crash occurs before all the individual writes complete, it is possible for some writes to complete while others do not. This can result in the data becoming unsynchronized. For mirrored volumes, it can cause two reads from the same region of the volume to return different results, if different mirrors are used to satisfy the read request. In the case of RAID-5 volumes, it can lead to parity corruption and incorrect data reconstruction.

VxVM ensures that all mirrors contain exactly the same data and that the data and parity in RAID-5 volumes agree. This process is called volume resynchronization. For volumes that are part of the disk group that is automatically imported at boot time (usually aliased as the reserved system-wide disk group, `bootdgc`), resynchronization takes place when the system reboots.

Not all volumes require resynchronization after a system failure. Volumes that were never written or that were quiescent (that is, had no active I/O) when the system failure occurred could not have had outstanding writes and do not require resynchronization.

Dirty flags

VxVM records when a volume is first written to and marks it as dirty. When a volume is closed by all processes or stopped cleanly by the administrator, and all writes have been completed, VxVM removes the dirty flag for the volume. Only volumes that are marked dirty require resynchronization.

Resynchronization process

The process of resynchronization depends on the type of volume. For mirrored volumes, resynchronization is done by placing the volume in recovery mode (also called read-writeback recovery mode). Resynchronization of data in the volume is done in the background. This allows the volume to be available for use while recovery is taking place. RAID-5 volumes that contain RAID-5 logs can “replay” those logs. If no logs are available, the volume is placed in reconstruct-recovery mode and all parity is regenerated.

Resynchronization can impact system performance. The recovery process reduces some of this impact by spreading the recoveries to avoid stressing a specific disk or controller.

For large volumes or for a large number of volumes, the resynchronization process can take time. These effects can be minimized by using dirty region logging (DRL) and FastResync (fast mirror resynchronization) for mirrored volumes, or by using RAID-5 logs for RAID-5 volumes.

See [“Dirty region logging”](#) on page 127.

For mirrored volumes used by Oracle, you can use the SmartSync feature, which further improves performance.

See [“SmartSync recovery accelerator”](#) on page 128.

Hot-relocation

Hot-relocation is a feature that allows a system to react automatically to I/O failures on redundant objects (mirrored or RAID-5 volumes) in VxVM and restore redundancy and access to those objects. VxVM detects I/O failures on objects and relocates the affected subdisks. The subdisks are relocated to disks designated as spare disks or to free space within the disk group. VxVM then reconstructs the objects that existed before the failure and makes them accessible again.

When a partial disk failure occurs (that is, a failure affecting only some subdisks on a disk), redundant data on the failed portion of the disk is relocated. Existing volumes on the unaffected portions of the disk remain accessible.

See [“How hot-relocation works”](#) on page 772.

Dirty region logging

Dirty region logging (DRL), if enabled, speeds recovery of mirrored volumes after a system crash. DRL tracks the regions that have changed due to I/O writes to a mirrored volume. DRL uses this information to recover only those portions of the volume.

If DRL is not used and a system failure occurs, all mirrors of the volumes must be restored to a consistent state. Restoration is done by copying the full contents of the volume between its mirrors. This process can be lengthy and I/O intensive.

Note: DRL adds a small I/O overhead for most write access patterns. This overhead is reduced by using SmartSync.

If an instant snap DCO volume is associated with a volume, a portion of the DCO volume can be used to store the DRL log. There is no need to create a separate DRL log for a volume which has an instant snap DCO volume.

Log subdisks and plexes

DRL log subdisks store the dirty region log of a mirrored volume that has DRL enabled. A volume with DRL has at least one log subdisk; multiple log subdisks can be used to mirror the dirty region log. Each log subdisk is associated with one plex of the volume. Only one log subdisk can exist per plex. If the plex contains only a log subdisk and no data subdisks, that plex is referred to as a log plex.

The log subdisk can also be associated with a regular plex that contains data subdisks. In that case, the log subdisk risks becoming unavailable if the plex must be detached due to the failure of one of its data subdisks.

If the `vxassist` command is used to create a dirty region log, it creates a log plex containing a single log subdisk by default. A dirty region log can also be set up manually by creating a log subdisk and associating it with a plex. The plex then contains both a log and data subdisks.

Sequential DRL

Some volumes, such as those that are used for database replay logs, are written sequentially and do not benefit from delayed cleaning of the DRL bits. For these volumes, sequential DRL can be used to limit the number of dirty regions. This allows for faster recovery. However, if applied to volumes that are written to randomly, sequential DRL can be a performance bottleneck as it limits the number of parallel writes that can be carried out.

The maximum number of dirty regions allowed for sequential DRL is controlled by a tunable as detailed in the description of `voldr1_max_seq_dirty`.

SmartSync recovery accelerator

The SmartSync feature of Veritas Volume Manager increases the availability of mirrored volumes by only resynchronizing changed data. (The process of resynchronizing mirrored databases is also sometimes referred to as resilvering.) SmartSync reduces the time required to restore consistency, freeing more I/O bandwidth for business-critical applications. SmartSync uses an extended interface between VxVM volumes, VxFS file systems, and the Oracle database to avoid unnecessary work during mirror resynchronization and to reduce the I/O overhead of the DRL. For example, Oracle® automatically takes advantage of SmartSync to perform database resynchronization when it is available.

Note: To use SmartSync with volumes that contain file systems, see the discussion of the Oracle Resilvering feature of Veritas File System (VxFS).

The following section describes how to configure VxVM raw volumes and SmartSync. The database uses the following types of volumes:

- Data volumes are the volumes used by the database (control files and tablespace files).
- Redo log volumes contain redo logs of the database.

SmartSync works with these two types of volumes differently, so they must be configured as described in the following sections.

To enable the use of SmartSync with database volumes in shared disk groups, set the value of the `volcvm_smartsync` tunable to 1.

Data volume configuration

The recovery takes place when the database software is started, not at system startup. This reduces the overall impact of recovery when the system reboots. Because the recovery is controlled by the database, the recovery time for the volume is the resilvering time for the database (that is, the time required to replay the redo logs).

Because the database keeps its own logs, it is not necessary for VxVM to do logging. Data volumes should be configured as mirrored volumes without dirty region logs. In addition to improving recovery time, this avoids any run-time I/O overhead due to DRL, and improves normal database write access.

Redo log volume configuration

A redo log is a log of changes to the database data. Because the database does not maintain changes to the redo logs, it cannot provide information about which sections require resilvering. Redo logs are also written sequentially, and since traditional dirty region logs are most useful with randomly-written data, they are of minimal use for reducing recovery time for redo logs. However, VxVM can reduce the number of dirty regions by modifying the behavior of its dirty region logging feature to take advantage of sequential access patterns. Sequential DRL decreases the amount of data needing recovery and reduces recovery time impact on the system.

The enhanced interfaces for redo logs allow the database software to inform VxVM when a volume is to be used as a redo log. This allows VxVM to modify the DRL behavior of the volume to take advantage of the access patterns. Since the improved

recovery time depends on dirty region logs, redo log volumes should be configured as mirrored volumes with sequential DRL.


See “[Sequential DRL](#)” on page 128.

Volume snapshots

Veritas Volume Manager provides the capability for taking an image of a volume at a given point in time. Such an image is referred to as a volume snapshot. Such snapshots should not be confused with file system snapshots, which are point-in-time images of a Veritas File System.

[Figure 4-30](#) shows how a snapshot volume represents a copy of an original volume at a given point in time.

Figure 4-30 Volume snapshot as a point-in-time image of a volume



Even though the contents of the original volume can change, the snapshot volume preserves the contents of the original volume as they existed at an earlier time.

The snapshot volume provides a stable and independent base for making backups of the contents of the original volume, or for other applications such as decision support. In the figure, the contents of the snapshot volume are eventually resynchronized with the original volume at a later point in time.

Another possibility is to use the snapshot volume to restore the contents of the original volume. This may be useful if the contents of the original volume have become corrupted in some way.

Warning: If you write to the snapshot volume, it may no longer be suitable for use in restoring the contents of the original volume.

One type of volume snapshot in VxVM is the third-mirror break-off type. This name comes from its implementation where a snapshot plex (or third mirror) is added to a mirrored volume. The contents of the snapshot plex are then synchronized from the original plexes of the volume. When this synchronization is complete, the snapshot plex can be detached as a snapshot volume for use in backup or decision support applications. At a later time, the snapshot plex can be reattached to the original volume, requiring a full resynchronization of the snapshot plex's contents.

The FastResync feature was introduced to track writes to the original volume. This tracking means that only a partial, and therefore much faster, resynchronization is required on reattaching the snapshot plex. In later releases, the snapshot model was enhanced to allow snapshot volumes to contain more than a single plex, reattachment of a subset of a snapshot volume's plexes, and persistence of FastResync across system reboots or cluster restarts.

Release 4.0 of VxVM introduced full-sized instant snapshots and space-optimized instant snapshots, which offer advantages over traditional third-mirror snapshots such as immediate availability and easier configuration and administration. You can also use the third-mirror break-off usage model with full-sized snapshots, where this is necessary for write-intensive applications.

See [“How full-sized instant snapshots work”](#) on page 587.

For information about how and when to use volume snapshots, see the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

See the `vxassist(1M)` manual page.

See the `vxsnap(1M)` manual page.

Comparison of snapshot features

[Table 4-1](#) compares the features of the various types of snapshots that are supported in VxVM.

Table 4-1 Comparison of snapshot features for supported snapshot types

Snapshot feature	Full-sized instant (vxsnap)	Space-optimized instant (vxsnap)	Break-off (vxassist or vxsnap)
Immediately available for use on creation	Yes	Yes	No

Table 4-1 Comparison of snapshot features for supported snapshot types
(continued)

Snapshot feature	Full-sized instant (vxsnap)	Space-optimized instant (vxsnap)	Break-off (vxassist or vxsnap)
Requires less storage space than original volume	No	Yes	No
Can be reattached to original volume	Yes	No	Yes
Can be used to restore contents of original volume	Yes	Yes	Yes
Can quickly be refreshed without being reattached	Yes	Yes	No
Snapshot hierarchy can be split	Yes	No	No
Can be moved into separate disk group from original volume	Yes	No	Yes
Can be turned into an independent volume	Yes	No	Yes
FastResync ability persists across system reboots or cluster restarts	Yes	Yes	Yes
Synchronization can be controlled	Yes	No	No
Can be moved off-host	Yes	No	Yes

Full-sized instant snapshots are easier to configure and offer more flexibility of use than do traditional third-mirror break-off snapshots. For preference, new volumes should be configured to use snapshots that have been created using the `vxsnap` command rather than using the `vxassist` command. Legacy volumes can also be reconfigured to use `vxsnap` snapshots, but this requires rewriting of administration scripts that assume the `vxassist` snapshot model.

FastResync

Note: Only certain Storage Foundation and High Availability Solutions products have a license to use this feature.

The FastResync feature (previously called Fast Mirror Resynchronization or FMR) performs quick and efficient resynchronization of stale mirrors (a mirror that is not synchronized). This feature increases the efficiency of the VxVM snapshot mechanism, and improves the performance of operations such as backup and decision support applications. Typically, these operations require that the volume is quiescent, and that they are not impeded by updates to the volume by other activities on the system. To achieve these goals, the snapshot mechanism in VxVM creates an exact copy of a primary volume at an instant in time. After a snapshot is taken, it can be accessed independently of the volume from which it was taken.

In a Cluster Volume Manager (CVM) environment with shared access to storage, it is possible to eliminate the resource contention and performance overhead of using a snapshot simply by accessing it from a different node.

How FastResync works

FastResync provides the following enhancements to VxVM:

Faster mirror resynchronization FastResync optimizes mirror resynchronization by keeping track of updates to stored data that have been missed by a mirror. (A mirror may be unavailable because it has been detached from its volume, either automatically by VxVM as the result of an error, or directly by an administrator using a utility such as `vxplex` or `vxassist`. A returning mirror is a mirror that was previously detached and is in the process of being re-attached to its original volume as the result of the `vxrecover` or `vxplex att` operation.) When a mirror returns to service, only the updates that it has missed need to be re-applied to resynchronize it. This requires much less effort than the traditional method of copying all the stored data to the returning mirror.

Once FastResync has been enabled on a volume, it does not alter how you administer mirrors. The only visible effect is that repair operations conclude more quickly.

Re-use of snapshots FastResync allows you to refresh and re-use snapshots rather than discard them. You can quickly re-associate (snap back) snapshot plexes with their original volumes. This reduces the system overhead required to perform cyclical operations such as backups that rely on the volume snapshots.

FastResync can be implemented in one of two ways:

Non-persistent FastResync Non-persistent FastResync allocates its change maps in memory. The maps do not reside on disk nor in persistent store.

See [“How non-persistent FastResync works with snapshots”](#) on page 134.

Persistent FastResync Persistent FastResync keeps the FastResync maps on disk so that they can survive system reboots, system crashes and cluster crashes.

See [“How persistent FastResync works with snapshots”](#) on page 135.

How non-persistent FastResync works with snapshots

If FastResync is enabled on a volume before a snapshot is taken, the snapshot feature of VxVM uses FastResync change tracking to record updates to the original volume after a snapshot plex is created. When the `snapback` option is used to reattach the snapshot plex, the changes that FastResync recorded are used to resynchronize the volume during the snapback. This behavior considerably reduces the time needed to resynchronize the volume.

Non-persistent FastResync uses a map in memory to implement change tracking. The map does not exist on disk or in persistent store. The advantage of non-persistent FastResync is that updates to the FastResync map have little impact on I/O performance, because no disk updates are performed. However, FastResync must remain enabled until the snapshot is reattached, and the system cannot be rebooted. If FastResync is disabled or the system is rebooted, the information in the map is lost and a full resynchronization is required on snapback.

This limitation can be overcome for volumes in cluster-shareable disk groups, provided that at least one of the nodes in the cluster remained running to preserve the FastResync map in its memory. However, a node crash in a High Availability (HA) environment requires the full resynchronization of a mirror when it is reattached to its parent volume.

Each bit in the FastResync map represents a contiguous number of blocks in a volume's address space. The default size of the map is 4 blocks. The kernel tunable `vol_fmr_logsz` can be used to limit the maximum size in blocks of the map

For information about tuning VxVM, see the *Veritas Storage Foundation and High Availability Tuning Guide*.

How persistent FastResync works with snapshots

Persistent FastResync keeps the FastResync maps on disk so that they can survive system reboots, system crashes and cluster crashes. Persistent FastResync uses a map in a data change object (DCO) volume on disk to implement change tracking. Each bit in the map represents a contiguous number of blocks in a volume's address space.

Persistent FastResync can also track the association between volumes and their snapshot volumes after they are moved into different disk groups. When the disk groups are rejoined, this allows the snapshot plexes to be quickly resynchronized. This ability is not supported by non-persistent FastResync.


See [“Reorganizing the contents of disk groups”](#) on page 814.

When persistent FastResync is enabled on a volume or on a snapshot volume, a data change object (DCO) and a DCO volume are associated with the volume.

See [“DCO volume versioning ”](#) on page 138.

[Figure 4-31](#) shows an example of a mirrored volume with two plexes on which persistent FastResync is enabled.

Figure 4-31 Mirrored volume with persistent FastResync enabled




Associated with the volume are a DCO object and a DCO volume with two plexes.

Create an instant snapshot by using the `vxsnap make` command, or create a traditional third-mirror snapshot by using the `vxassist snapstart` command.

Figure 4-32 shows how a snapshot plex is set up in the volume, and how a disabled DCO plex is associated with it.

Figure 4-32 Mirrored volume after completion of a snapstart operation




Multiple snapshot plexes and associated DCO plexes may be created in the volume by re-running the `vxassist snapstart` command for traditional snapshots, or the `vxsnap make` command for space-optimized snapshots. You can create up to a total of 32 plexes (data and log) in a volume.

A traditional snapshot volume is created from a snapshot plex by running the `vxassist snapshot` operation on the volume. For instant snapshots, however, the `vxsnap make` command makes an instant snapshot volume immediately available for use. There is no need to run an additional command.

Figure 4-33 shows how the creation of the snapshot volume also sets up a DCO object and a DCO volume for the snapshot volume.

Figure 4-33 Mirrored volume and snapshot volume after completion of a snapshot operation



The DCO volume contains the single DCO plex that was associated with the snapshot plex. If two snapshot plexes were taken to form the snapshot volume, the DCO volume would contain two plexes. For space-optimized instant snapshots, the DCO object and DCO volume are associated with a snapshot volume that is created on a cache object and not on a VM disk.

Associated with both the original volume and the snapshot volume are snap objects. The snap object for the original volume points to the snapshot volume, and the snap object for the snapshot volume points to the original volume. This allows VxVM to track the relationship between volumes and their snapshots even if they are moved into different disk groups.

The snap objects in the original volume and snapshot volume are automatically deleted in the following circumstances:

- For traditional snapshots, the `vxassist snapback` operation is run to return all of the plexes of the snapshot volume to the original volume.
- For traditional snapshots, the `vxassist snapclear` operation is run on a volume to break the association between the original volume and the snapshot volume.

If the volumes are in different disk groups, the command must be run separately on each volume.

- For full-sized instant snapshots, the `vxsnap reattach` operation is run to return all of the plexes of the snapshot volume to the original volume.
- For full-sized instant snapshots, the `vxsnap dis` or `vxsnap split` operations are run on a volume to break the association between the original volume and the snapshot volume. If the volumes are in different disk groups, the command must be run separately on each volume.

Note: The `vxsnap reattach`, `dis` and `split` operations are not supported for space-optimized instant snapshots.

See the `vxassist`(1M) manual page.

See the `vxsnap`(1M) manual page.

DCO volume versioning

Persistent FastResync uses a data change object (DCO) and a DCO volume to hold the FastResync maps.

This release of Veritas Volume Manager (VxVM) supports the following DCO volume versions:

Instant snap DCO volume layout	Previously known as Version 20 DCO volume layout, this version of the DCO layout supports instant snapshots of volumes. This type of DCO manages the FastResync maps, and also manages DRL recovery maps and special maps called copymaps that allow instant snapshot operations to resume correctly following a system crash.
Version 0 DCO volume layout	This version of the DCO volume layout only supports legacy snapshots (<code>vxassist</code> snapshots). The DCO object manages information about the FastResync maps. These maps track writes to the original volume and to each of up to 32 snapshot volumes since the last <code>snapshot</code> operation. Each plex of the DCO volume on disk holds 33 maps, each of which is 4 blocks in size by default. VxVM software continues to support the version 0 (zero) layout for legacy volumes.

Instant snap (version 20) DCO volume layout

The instant snap data change object (DCO) supports full-sized and space-optimized instant snapshots. Traditional third-mirror volume snapshots that are administered using the `vxassist` command are not supported with this DCO layout.

Introduced in Veritas Volume Manager (VxVM) 4.0, the instant snap DCO volume layout is also known as a version 20 DCO volume layout. This type of DCO is used not only to manage the FastResync maps, but also to manage DRL recovery maps and special maps called copymaps that allow instant snapshot operations to resume correctly following a system crash.

See “[Dirty region logging](#)” on page 127.

Each bit in a map represents a region (a contiguous number of blocks) in a volume’s address space. A region represents the smallest portion of a volume for which changes are recorded in a map. A write to a single byte of storage anywhere within a region is treated in the same way as a write to the entire region.

In Storage Foundation Cluster File System High Availability 6.0, the volume layout of an instant snap DCO has been changed to improve the I/O performance and scalability of instant snapshots. The change in layout does not alter how you administer instant snapshots. The only visible affect is in improved I/O performance and in some cases, increased size of DCO volume.

The layout of an instant snap DCO volume uses dynamic creation of maps on the preallocated storage. The size of the DRL (Dirty region logging) map does not depend on volume size. You can configure the size of the DRL by using the option `drlmapsz` while creating the DCO volume. By default, the size of the DRL is 1MB.

For CVM configurations, each node has a dedicated DRL map which gets allocated during first write on that node. By default, the size of DCO volume accommodates 32 DRL maps, an accumulator, and 16 per-volume maps (includes a DRL recovery map, a detach map to track detached plexes and remaining 14 maps for tracking snapshots).

The size of the DCO plex can be estimated using the following formula:

$$DCO\_volume\_size = (32 * drlmapsize + acmsize + 16 * per\_volume\_map\_size)$$

where:

$$acmsize = (volume\_size / (region\_size * 4))$$
$$per\_volume\_map\_size = (volume\_size / region\_size * 8)$$
$$drlmapsize = 1M, \text{ by default}$$

For a 100GB volume, the size of the DCO volume with the default *regionsize* of 64KB is approximately 36MB.

Create the DCOs for instant snapshots by using the `vxsnap prepare` command or by specifying the options `logtype=dco dconversion=20` while creating a volume with the `vxassist make` command.

Version 0 DCO volume layout

The version 0 DCO volume layout supports only traditional (third-mirror) volume snapshots that are administered using the `vxassist` command. Full-sized and space-optimized instant snapshots are not supported with this DCO layout.

The size of each map can be changed by specifying the `dcolen` attribute to the `vxassist` command when the volume is created. The default value of `dcolen` is 132 blocks (the plex contains 33 maps, each of length 4 blocks). To use a larger map size, multiply the desired map size by 33 to calculate the value of `dcolen`. For example, to use an 8-block map, specify `dcolen=264`. The maximum possible map size is 64 blocks, which corresponds to a `dcolen` value of 2112 blocks.

The size of a DCO plex is rounded up to the nearest integer multiple of the disk group alignment value. The alignment value is 8KB for disk groups that support the Cross-platform Data Sharing (CDS) feature. Otherwise, the alignment value is 1 block.

Effect of growing a volume on the FastResync map

It is possible to grow the replica volume, or the original volume, and still use FastResync. According to the DCO volume layout, growing the volume has the following different effects on the map that FastResync uses to track changes to the original volume:

- For an instant snap DCO volume, the size of the map is increased and the size of the region that is tracked by each bit in the map stays the same.
- For a version 0 DCO volume, the size of the map remains the same and the region size is increased.

In either case, the part of the map that corresponds to the grown area of the volume is marked as “dirty” so that this area is resynchronized. The `snapback` operation fails if it attempts to create an incomplete snapshot plex. In such cases, you must grow the replica volume, or the original volume, before invoking any of the commands `vxsnap reattach`, `vxsnap restore`, or `vxassist snapback`. Growing the two volumes separately can lead to a snapshot that shares physical disks with another mirror in the volume. To prevent this, grow the volume after the `snapback` command is complete.

FastResync limitations

The following limitations apply to FastResync:

- Persistent FastResync is supported for RAID-5 volumes, but this prevents the use of the relayout or resize operations on the volume while a DCO is associated with it.
- Neither non-persistent nor persistent FastResync can be used to resynchronize mirrors after a system crash. Dirty region logging (DRL), which can coexist with FastResync, should be used for this purpose. In VxVM 4.0 and later releases, DRL logs may be stored in an instant snap DCO volume.
- When a subdisk is relocated, the entire plex is marked “dirty” and a full resynchronization becomes necessary.
- If a snapshot volume is split off into another disk group, non-persistent FastResync cannot be used to resynchronize the snapshot plexes with the original volume when the disk group is rejoined with the original volume’s disk group. Persistent FastResync must be used for this purpose.
- If you move or split an original volume (on which persistent FastResync is enabled) into another disk group, and then move or join it to a snapshot volume’s disk group, you cannot use `vxassist snapback` to resynchronize traditional snapshot plexes with the original volume. This restriction arises because a snapshot volume references the original volume by its record ID at the time that the snapshot volume was created. Moving the original volume to a different disk group changes the volume’s record ID, and so breaks the association. However, in such a case, you can use the `vxplex snapback` command with the `-f` (force) option to perform the snapback.

Note: This restriction only applies to traditional snapshots. It does not apply to instant snapshots.

- Any operation that changes the layout of a replica volume can mark the FastResync change map for that snapshot “dirty” and require a full resynchronization during snapback. Operations that cause this include subdisk split, subdisk move, and online relayout of the replica. It is safe to perform these operations after the snapshot is completed.

See the `vxassist` (1M) manual page.

See the `vxplex` (1M) manual page.

See the `vxvol` (1M) manual page.

Volume sets

Volume sets are an enhancement to VxVM that allow several volumes to be represented by a single logical object. All I/O from and to the underlying volumes is directed via the I/O interfaces of the volume set. The Veritas File System (VxFS) uses volume sets to manage multi-volume file systems and the SmartTier feature. This feature allows VxFS to make best use of the different performance and availability characteristics of the underlying volumes. For example, file system metadata can be stored on volumes with higher redundancy, and user data on volumes with better performance.


See [“Creating a volume set”](#) on page 702.

Configuration of volumes on SAN storage

Storage Area Networks (SANs) provide a networking paradigm that provides easily reconfigurable connectivity between any subset of computers, disk storage, and switches. A SAN can contain a huge number of devices connected using switched fabric. A SAN that has thousands or tens of thousands of connected devices is difficult to administer using a simple disk group model. Veritas CommandCentral Storage software allows you to configure storage groups and storage accounts. Using the CommandCentral Storage software, you can allocate SAN storage more prudently and administer your complex SAN environments more effectively.

Note: This feature of `vxassist` is designed to work in conjunction with SAL (SAN Access Layer) in Veritas CommandCentral Storage. When VxVM with SAN-aware `vxassist` is installed on a host where SAL is also installed, it is recommended that you create a user named `root` under SAL. This allows `vxassist` to use the `root` login to contact the SAL daemon (`salld`) on the primary SAL server without needing to specify the `sal_username` attribute to `vxassist`.

[Figure 4-34](#), shows how you might choose to set up storage groups within a SAN.

Figure 4-34 Dividing a Storage Area Network into storage groups

In this example, the boundaries of the storage groups are based on the performance characteristics of different makes of disk array and on geographic location.

The `vxassist` utility in Veritas Volume Manager understands storage groups that you have defined using the CommandCentral Storage software. `vxassist` supports a simple language that you can use to specify how disks are to be allocated from pre-defined storage groups. This specification language defines the confinement and separation criteria that `vxassist` applies to the available storage to choose disks for creating, resizing or moving a volume.

To use the CommandCentral Storage storage groups with `vxassist`, perform the following steps in the order listed:

- Use the CommandCentral Storage software to define one or more storage groups. Note that zoning is not an issue as it is completely independent of storage group creation.
- Use the CommandCentral Storage software to attach attribute-value pairs to each storage group's property sheet. Typically, you would assign values for the following attributes: location, storage group, and protection.
- Use the `vxspcshow` command to discover the device names of disks that have a specified set of attributes, or to list the attributes of specified disks.

- Use the `vxdiskadm` command to configure the disks that you found in the previous step into VxVM disk groups.
- Use `vxassist` to create volumes on disks that are selected by matching specified criteria for the values of storage group attributes. The usual restriction applies that a volume may only be created using disks from a single disk group.

How CVM works

This chapter includes the following topics:

- [About the cluster functionality of VxVM](#)
- [Overview of clustering](#)
- [CVM initialization and configuration](#)
- [Dirty region logging in cluster environments](#)
- [Multiple host failover configurations](#)

About the cluster functionality of VxVM

A cluster consists of a number of hosts or nodes that share a set of disks. The following are the main benefits of cluster configurations:

Availability

If one node fails, the other nodes can still access the shared disks. When configured with suitable software, mission-critical applications can continue running by transferring their execution to a standby node in the cluster. This ability to provide continuous uninterrupted service by switching to redundant hardware is commonly termed failover.

Failover is transparent to users and high-level applications for database and file-sharing. You must configure cluster management software, such as Veritas Cluster Server (VCS), to monitor systems and services, and to restart applications on another node in the event of either hardware or software failure. VCS also allows you to perform general administration tasks such as making nodes join or leave a cluster.

Note that a standby node need not remain idle. It could be used to serve other applications in parallel.

Off-host
processing

Clusters can reduce contention for system resources by performing activities such as backup, decision support and report generation on the more lightly loaded nodes of the cluster. This allows businesses to derive enhanced value from their investment in cluster systems.

The nodes can simultaneously access and manage a set of disks or LUNs under VxVM control. The same logical view of disk configuration and any changes to this view are available on all the nodes. When the CVM functionality is enabled, all cluster nodes can share VxVM objects such as shared disk groups. Private disk groups are supported in the same way as in a non-clustered environment. This chapter discusses the cluster functionality that is provided with VxVM.

Products such as Veritas Storage Foundation Cluster File System (SFCFS), and Veritas Cluster Server (VCS) are separately licensed, and are not included with Veritas Volume Manager. See the documentation provided with those products for more information about them.

Veritas Dynamic Multi-Pathing (DMP) can be used in a clustered environment.

See [“DMP in a clustered environment”](#) on page 180.

Campus cluster configurations (also known as stretch cluster or remote mirror configurations) can also be configured and administered.

See [“About sites and remote mirrors”](#) on page 464.

Overview of clustering

Tightly-coupled cluster systems are common in the realm of enterprise-scale mission-critical data processing. The primary advantage of clusters is protection against hardware failure. Should the primary node fail or otherwise become unavailable, applications can continue to run by transferring their execution to standby nodes in the cluster. This ability to provide continuous availability of service by switching to redundant hardware is commonly termed failover.

Another major advantage of clustered systems is their ability to reduce contention for system resources caused by activities such as backup, decision support and report generation. Businesses can derive enhanced value from their investment in cluster systems by performing such operations on lightly loaded nodes in the cluster rather than on the heavily loaded nodes that answer requests for service. This ability to perform some operations on the lightly loaded nodes is commonly termed load balancing.


Overview of cluster volume management

Over the past several years, parallel applications using shared data access have become increasingly popular. Examples of commercially available applications include Oracle Real Application Clusters™ (RAC), Sybase Adaptive Server®, and Informatica Enterprise Cluster Edition. In addition, the semantics of Network File System (NFS), File Transfer Protocol (FTP), and Network News Transfer Protocol (NNTP) allow these workloads to be served by shared data access clusters. Finally, numerous organizations have developed internal applications that take advantage of shared data access clusters.

The cluster functionality of VxVM (CVM) works together with the cluster monitor daemon that is provided by VCS or by the host operating system. The cluster monitor informs VxVM of changes in cluster membership. Each node starts up independently and has its own cluster monitor plus its own copies of the operating system and VxVM/CVM. When a node joins a cluster, it gains access to shared disk groups and volumes. When a node leaves a cluster, it loses access to these shared objects. A node joins a cluster when you issue the appropriate command on that node.

Warning: The CVM functionality of VxVM is supported only when used in conjunction with a cluster monitor that has been configured correctly to work with VxVM.

[Figure 5-1](#) shows a simple cluster arrangement consisting of four nodes with similar or identical hardware characteristics (CPUs, RAM and host adapters), and configured with identical software (including the operating system).

Figure 5-1 Example of a 4-node CVM cluster

To the cluster monitor, all nodes are the same. VxVM objects configured within shared disk groups can potentially be accessed by all nodes that join the cluster. However, the CVM functionality of VxVM requires that one node act as the master node; all other nodes in the cluster are slave nodes. Any node is capable of being the master node, and it is responsible for coordinating certain VxVM activities.

In this example, node 0 is configured as the CVM master node and nodes 1, 2 and 3 are configured as CVM slave nodes. The nodes are fully connected by a private network and they are also separately connected to shared external storage (either disk arrays or JBODs: just a bunch of disks) via SCSI or Fibre Channel in a Storage Area Network (SAN).

In this example, each node has two independent paths to the disks, which are configured in one or more cluster-shareable disk groups. Multiple paths provide resilience against failure of one of the paths, but this is not a requirement for cluster configuration. Disks may also be connected by single paths.

The private network allows the nodes to share information about system resources and about each other's state. Using the private network, any node can recognize which other nodes are currently active, which are joining or leaving the cluster, and which have failed. The private network requires at least two communication channels to provide redundancy against one of the channels failing. If only one channel were used, its failure would be indistinguishable from node failure—a condition known as network partitioning.

You can run commands that configure or reconfigure VxVM objects on any node in the cluster. These tasks include setting up shared disk groups, creating and reconfiguring volumes, and performing snapshot operations.

The first node to join a cluster performs the function of master node. If the master node leaves a cluster, one of the slave nodes is chosen to be the new master.

See [“Methods to control CVM master selection”](#) on page 367.

Private and shared disk groups

The following types of disk groups are defined:

Private disk group Belongs to only one node. A private disk group can only be imported by one system. LUNs in a private disk group may be physically accessible from one or more systems, but access is restricted to only one system at a time.

The boot disk group (usually aliased by the reserved disk group name `bootdg`) is always a private disk group.

Shared disk group Can be shared by all nodes. A shared (or cluster-shareable) disk group is imported by all cluster nodes. LUNs in a shared disk group must be physically accessible from all systems that may join the cluster.

In a CVM cluster, most disk groups are shared. LUNs in a shared disk group are accessible from all nodes in a cluster, allowing applications on multiple cluster nodes to simultaneously access the same LUN. A volume in a shared disk group can be simultaneously accessed by more than one node in the cluster, subject to license key and disk group activation mode restrictions.

You can use the `vxchg` command to designate a disk group as cluster-shareable.

See [“Importing disk groups as shared”](#) on page 382.

When a disk group is imported as cluster-shareable for one node, each disk header is marked with the cluster ID. As each node subsequently joins the cluster, it recognizes the disk group as being cluster-shareable and imports it. In contrast, a private disk group's disk headers are marked with the individual node's host name. As system administrator, you can import or deport a shared disk group at any time; the operation takes place in a distributed fashion on all nodes.

Each LUN is marked with a unique disk ID. When cluster functionality for VxVM starts on the master, it imports all shared disk groups (except for any that do not have the `autoimport` attribute set). When a slave tries to join a cluster, the master sends it a list of the disk IDs that it has imported, and the slave checks to see if it can access them all. If the slave cannot access one of the listed disks, it abandons its attempt to join the cluster. If it can access all of the listed disks, it joins the cluster

and imports the same shared disk groups as the master. When a node leaves the cluster gracefully, it departs all its imported shared disk groups, but they remain imported on the surviving nodes.

Reconfiguring a shared disk group is performed with the cooperation of all nodes. Configuration changes to the disk group are initiated by the master, and happen simultaneously on all nodes and the changes are identical. Such changes are atomic in nature, which means that they either occur simultaneously on all nodes or not at all.

Whether all members of the cluster have simultaneous read and write access to a cluster-shareable disk group depends on its activation mode setting.

See [“Activation modes of shared disk groups”](#) on page 150.

The data contained in a cluster-shareable disk group is available as long as at least one node is active in the cluster. The failure of a cluster node does not affect access by the remaining active nodes. Regardless of which node accesses a cluster-shareable disk group, the configuration of the disk group looks the same.

Warning: Applications running on each node can access the data on the VM disks simultaneously. VxVM does not protect against simultaneous writes to shared volumes by more than one node. It is assumed that applications control consistency (by using Veritas Cluster File System or a distributed lock manager, for example).

Activation modes of shared disk groups

A shared disk group must be activated on a node in order for the volumes in the disk group to become accessible for application I/O from that node. The ability of applications to read from or to write to volumes is dictated by the activation mode of a shared disk group. Valid activation modes for a shared disk group are `exclusivewrite`, `readonly`, `sharedread`, `sharedwrite`, and `off` (inactive).

The default activation mode for shared disk groups is `sharedwrite`.

Special uses of clusters, such as high availability (HA) applications and off-host backup, can use disk group activation to explicitly control volume access from different nodes in the cluster

[Table 5-1](#) describes the activation modes.

Table 5-1 Activation modes for shared disk groups

Activation mode	Description
<code>exclusivewrite</code> (<code>ew</code>)	The node has exclusive write access to the disk group. No other node can activate the disk group for write access.

Table 5-1 Activation modes for shared disk groups (*continued*)

Activation mode	Description
<code>readonly (ro)</code>	The node has read access to the disk group and denies write access for all other nodes in the cluster. The node has no write access to the disk group. Attempts to activate a disk group for either of the write modes on other nodes fail.
<code>sharedread (sr)</code>	The node has read access to the disk group. The node has no write access to the disk group, however other nodes can obtain write access.
<code>sharedwrite (sw)</code>	The node has write access to the disk group. Attempts to activate the disk group for shared read and shared write access succeed. Attempts to activate the disk group for exclusive write and read-only access fail.
<code>off</code>	The node has neither read nor write access to the disk group. Query operations on the disk group are permitted.

Table 5-2 summarizes the allowed and conflicting activation modes for shared disk groups.

Table 5-2 Allowed and conflicting activation modes

Disk group activated in cluster as...	Attempt to activate disk group on another node as...			
	exclusive-write	readonly	sharedread	sharedwrite
<code>exclusivewrite</code>	Fails	Fails	Succeeds	Fails
<code>readonly</code>	Fails	Succeeds	Succeeds	Fails
<code>sharedread</code>	Succeeds	Succeeds	Succeeds	Succeeds
<code>sharedwrite</code>	Fails	Fails	Succeeds	Succeeds

Shared disk groups can be automatically activated in a specified mode when the disk group is created or imported. To control automatic activation of shared disk groups, create a defaults file `/etc/default/vxdg` containing the following lines:

```
enable_activation=true
default_activation_mode=activation-mode
```

The *activation-mode* is one of `exclusivewrite`, `readonly`, `sharedread`, `sharedwrite`, or `off`.

When a shared disk group is created or imported, it is activated in the specified mode. When a node joins the cluster, all shared disk groups accessible from the node are activated in the specified mode.

The activation mode of a disk group controls volume I/O from different nodes in the cluster. It is not possible to activate a disk group on a given node if it is activated in a conflicting mode on another node in the cluster. When enabling activation using the defaults file, it is recommended that the file be consistent on all nodes in the cluster as in [Table 5-2](#). Otherwise, the results of activation are unpredictable.

If the defaults file is edited while the `vxconfigd` daemon is already running, run the `/sbin/vxconfigd -k -x syslog` command on all nodes to restart the process.

If the default activation mode is anything other than `off`, an activation following a cluster join, or a disk group creation or import can fail if another node in the cluster has activated the disk group in a conflicting mode.

To display the activation mode for a shared disk group, use the `vx dg list diskgroup` command.

See [“Listing shared disk groups”](#) on page 380.

You can also use the `vx dg` command to change the activation mode on a shared disk group.

See [“Changing the activation mode on a shared disk group”](#) on page 384.

It is also possible to configure a volume so that it can only be opened by a single node in a cluster.

See [“Creating volumes with exclusive open access by a node”](#) on page 385.

See [“Setting exclusive open access to a volume by a node”](#) on page 385.

CVM resiliency to storage connectivity failures

A shared disk group provides concurrent read and write access to the volumes that it contains for all nodes in a cluster. A shared disk group can be created on any node of the cluster.

CVM maintains configuration copies on the shared disks. The configuration of a shared disk group is on the private region of the disk. CVM maintains the configuration copies on shared storage.

Cluster Volume Manager (CVM) manages the configuration of a shared disk group as follows:

- All nodes in the cluster see the same storage configuration.

- Commands to change the configuration are sent to the master node.
- The master node performs changes to the configuration and propagates the changes to the slave nodes in the cluster.

A shared disk group has the same configuration on all nodes. That is, all nodes have the same view of the storage. Some error scenarios create exceptions to this guideline. If a node loses connectivity to the shared storage, CVM handles the storage disconnectivity with the appropriate error handling.

CVM responds to storage disconnectivity in the following ways:

- Availability to data disks.
The disk detach policy specifies how CVM handles connectivity errors for data disks, and whether plex availability or node availability should take precedence. You can specify one of the following values.
 - The global detach policy indicates that a local fault on one node detaches the plex on all nodes in the cluster.
 - The local detach policy indicates that a local fault has a local impact. This policy indicates that plex availability is given higher priority than the number of nodes that can access the plex.

See [“Disk detach policies”](#) on page 154.

- Availability to shared disk group configurations.
CVM maintains configuration copies on shared disks. The CVM master node is responsible for writing any change requests to the configuration copy to the disks. If the master loses access to a configuration copy, CVM redirects the write requests over the network to another node that has connectivity to the configuration copy. This behavior ensures that the disk group stays available. If the disk group version is less than 170, CVM handles disconnectivity according to the disk group failure policy (`dgfail_policy`).

See [“Availability of shared disk group configuration copies”](#) on page 153.

Availability of shared disk group configuration copies

CVM maintains connectivity to shared disk groups as long as at least one node in the cluster has access to the configuration copies. The master node performs changes to the configuration and propagates the changes to all of the slave nodes. If the master node loses access to the copies of the configuration, the master node sends the writes for the configuration change to a slave node that has access. The slave node implements the changes. This behavior enables the disk group to stay active, even when the master node loses access to the shared disk group configuration. If all nodes lose access to the shared disk group, the disk group is disabled.

In previous releases, you could configure a disk group failure policy for a shared disk group. In this release, the disk group failure policy is no longer supported for disk groups with the latest disk group version and cluster protocol version.

If the cluster protocol version is less than 110, or the disk group version is less than 170, the disk group failure policy (dgfail_policy) determines the behavior after failure.

See the documentation from the relevant release of VxVM.

Disk detach policies

The disk detach policy determines how CVM handles failures in the storage or connectivity to the storage.

For small mirrored volumes, non-mirrored volumes, volumes that use hardware mirrors, and volumes in private disk groups, there is no benefit in configuring the local detach policy. In most cases, Symantec recommends that you use the default global detach policy.

The following disk detach policies are available:

- Global detach policy
See [“Global detach policy”](#) on page 154.
- Local detach policy
See [“Local detach policy”](#) on page 155.

Global detach policy

The global detach policy specifies that on any I/O error, the plex seeing the error is detached cluster-wide for the volume. This behavior assures a symmetric access to the volumes on all nodes in the cluster. The global detach policy is the traditional and default policy for all nodes in the configuration.

The benefit with this policy is that the volume is still available for I/O on all nodes. If there is a read or write I/O failure on a slave node, the master node performs the usual I/O recovery operations to repair the failure. If required, the plex is detached from the volume, for the entire cluster. All nodes remain in the cluster and continue to perform I/O, but the redundancy of the mirrors is reduced.

The disadvantage is that redundancy is lost, because of the detached plex. Because one or more nodes in the cluster lose connectivity to a plex, the entire cluster loses access to that plex. This behavior means that a local fault on one node has a global impact on all nodes in the cluster.

The global detach policy also requires the overhead of reattaching the plex. When the problem that caused the I/O failure has been corrected, the disks should be re-attached. The mirrors that were detached must be recovered before the redundancy of the data can be restored.

In this release, if a node experiences failure to all the plexes of a mirrored volume, the volume is disabled locally. This prevents the behavior from previous releases wherein each plex was detached one after other and the volume was disabled globally.

Local detach policy

The local detach policy indicates that when a node has an I/O failure to a plex due to a local connectivity issue, the volume is disabled locally. The plex is not detached for the whole cluster. This behavior ensures that all of the plexes are available for I/O on the other nodes. Only the node or nodes that had the failure are affected.

The benefit with this policy is that the redundancy of the volume is protected, as much as possible. The local detach policy supports failover applications in large clusters where the redundancy of the volume is more important than the number of nodes that can access the volume. That is, you would prefer one or more nodes to lose write access to the volume, rather than lose a plex of the volume for all of the nodes in the cluster. This scenario typically applies when the volumes are mirrored, and a parallel application is used that can seamlessly provide the same service from the other nodes. For example, this option is not appropriate for fast failover configurations.

If there is a write failure on any node, CVM contacts all the nodes to see if the plex is still accessible to them. The write failure is considered local if at least one node still has access to the plex. For local failure, the volume is locally disabled.

If the write failure is due to a connectivity issue across all nodes of the cluster, the plex is detached.

Use the `vxchg` command to set the disk detach policy on a shared disk group.

Guidelines for choosing detach policies

In most cases it is recommended that you use the global detach policy, and particularly if any of the following conditions apply:

- If only non-mirrored, small mirrored, or hardware mirrored volumes are configured. This avoids the system overhead of the extra messaging that is required by the local detach policy.

The local detach policy may be suitable in the following cases:

- When large mirrored volumes are configured. Resynchronizing a reattached plex can degrade system performance. The local detach policy can avoid the need to detach the plex at all. (Alternatively, the dirty region logging (DRL) feature can be used to reduce the amount of resynchronization that is required.)
- For clusters with more than four nodes. Keeping an application running on a particular node is less critical when there are many nodes in a cluster. It may

be possible to configure the cluster management software to move an application to a node that has access to the volumes. In addition, load balancing may be able to move applications to a different volume from the one that experienced the I/O problem. This preserves data redundancy, and other nodes may still be able to perform I/O from/to the volumes on the disk.

[Table 5-3](#) compares the behavior of CVM when I/O failures occur under the different detach policies.

Table 5-3 Cluster behavior under I/O failure to a mirrored volume for disk detach policies

Nodes seeing the failure	Type of failure	Local detach policy	Global detach policy
One or more nodes, but not all. (Local failure)	I/O failure on one or more plexes in a volume (but not all).	Locally disables the volume. I/O fails to the volume from the nodes that cannot access the plex. I/O to the volume continues from other nodes.	Detaches the plex. I/O continues to other plexes in the volume.
All nodes. (Global failure)	I/O failure on one or more plexes in a volume (but not all).	Detaches the plex. No nodes can access the plex, so the plex is detached. I/O continues to other plexes in the volume.	Detaches the plex. No nodes can access the plex, so the plex is detached. I/O continues to other plexes in the volume.
One or more nodes, but not all. (Local failure)	I/O failure to all plexes in a volume.	Locally disables the volume. I/O fails to the volume from the nodes that cannot access the plex. I/O continues to the volume from other nodes.	I/O fails on the volume. The node cannot access any plex in the volume, so the I/O to the volume fails.

Table 5-3 Cluster behavior under I/O failure to a mirrored volume for disk detach policies (*continued*)

Nodes seeing the failure	Type of failure	Local detach policy	Global detach policy
All nodes. (Global failure)	I/O failure to all plexes in a volume.	Disables the volume. Since no plexes are available, the volume is not available for any I/Os. No plexes are detached.	Disables the volume. Since no plexes are available, the volume is not available for any I/Os. No plexes are detached.

Effect of disk connectivity on cluster reconfiguration

The detach policy does not change the requirement that a node joining a cluster must have access to all the disks in all shared disk groups. Similarly, a node that is removed from the cluster because of an I/O failure cannot rejoin the cluster until this requirement is met.

Limitations of shared disk groups

Only raw device access may be performed via CVM. It does not support shared access to file systems in shared volumes unless the appropriate software, such as Veritas Cluster File System, is installed and configured.

Note: The boot disk group (usually aliased as `bootdg`) cannot be made cluster-shareable. It must be private.

The cluster functionality of VxVM does not support RAID-5 volumes, or task monitoring for cluster-shareable disk groups. These features can, however, be used in private disk groups that are attached to specific nodes of a cluster or can be failed over to other nodes in the cluster.

If you have RAID-5 volumes in a private disk group that you wish to make shareable, you must first relayout the volumes as a supported volume type such as `stripe-mirror` or `mirror-stripe`. Online relayout of shared volumes is supported provided that it does not involve RAID-5 volumes.

If a shared disk group contains RAID-5 volumes, deport it and then reimport the disk group as private on one of the cluster nodes. Reorganize the volumes into

layouts that are supported for shared disk groups, and then deport and reimport the disk group as shared.

CVM initialization and configuration

Before any nodes can join a new cluster for the first time, you must supply certain configuration information during cluster monitor setup. This information is normally stored in some form of cluster monitor configuration database. The precise content and format of this information depends on the characteristics of the cluster monitor. The information required by VxVM is as follows:

- Cluster ID
- Node IDs
- Network addresses of nodes
- Port addresses

When a node joins the cluster, this information is automatically loaded into VxVM on that node at node startup time.

Note: The CVM functionality of VxVM is supported only when used with a cluster monitor that has been configured correctly to work with VxVM.

Use a cluster monitor such as GAB (Group Membership and Atomic Broadcast) in Veritas Cluster Service (VCS). For VCS, the Veritas product installer collects the required information to configure the cluster monitor.

The cluster monitor startup procedure effects node initialization, and brings up the various cluster components (such as VxVM with cluster support, the cluster monitor, and a distributed lock manager) on the node. Once this is complete, applications may be started. The cluster monitor startup procedure must be invoked on each node to be joined to the cluster.

For VxVM in a cluster environment, initialization consists of loading the cluster configuration information and joining the nodes in the cluster. The first node to join becomes the master node, and later nodes (slaves) join to the master. If two nodes join simultaneously, VxVM chooses the master. After a given node joins, that node has access to the shared disk groups and volumes.

Cluster reconfiguration

Cluster reconfiguration occurs if a node leaves or joins a cluster. Each node's cluster monitor continuously watches the other cluster nodes. When the membership of

the cluster changes, the cluster monitor informs VxVM for it to take appropriate action.

During cluster reconfiguration, VxVM suspends I/O to shared disks. I/O resumes when the reconfiguration completes. Applications may appear to freeze for a short time during reconfiguration.

If other operations, such as VxVM operations or recoveries, are in progress, cluster reconfiguration can be delayed until those operations complete. Volume reconfigurations do not take place at the same time as cluster reconfigurations. Depending on the circumstances, an operation may be held up and restarted later. In most cases, cluster reconfiguration takes precedence. However, if the volume reconfiguration is in the commit stage, it completes first.

See [“Volume reconfiguration”](#) on page 161.

See [“vxclustadm utility”](#) on page 159.

vxclustadm utility

The `vxclustadm` command provides an interface to the CVM functionality of VxVM when VCS is used as the cluster monitor. It is also called during cluster startup and shutdown. In the absence of a cluster monitor, `vxclustadm` can also be used to activate or deactivate the CVM functionality of VxVM on any node in a cluster.

The `startnode` keyword to `vxclustadm` starts CVM functionality on a cluster node by passing cluster configuration information to the VxVM kernel. In response to this command, the kernel and the VxVM configuration daemon, `vxconfigd`, perform initialization.

The `stopnode` keyword stops CVM functionality on a node. It waits for all outstanding I/O to complete and for all applications to close shared volumes.

The `setmaster` keyword migrates the CVM master to the specified node. The migration is an online operation. Symantec recommends that you switch the master when the cluster is not handling VxVM configuration changes or cluster reconfiguration operations.

The `reinit` keyword allows nodes to be added to or removed from a cluster without stopping the cluster. Before running this command, the cluster configuration file must have been updated with information about the supported nodes in the cluster.

The `nidmap` keyword prints a table showing the mapping between CVM node IDs in VxVM's cluster-support subsystem and node IDs in the cluster monitor. It also prints the state of the nodes in the cluster.

The `nodestate` keyword reports the state of a cluster node and also the reason for the last abort of the node as shown in this example:

```
# vxclustadm nodestate
```

```
state: out of cluster
reason: user initiated stop
```

Table 5-4 lists the various reasons that may be given for a node abort.

Table 5-4 Node abort messages

Reason	Description
cannot find disk on slave node	Missing disk or bad disk on the slave node.
cannot obtain configuration data	The node cannot read the configuration data due to an error such as disk failure.
cluster device open failed	Open of a cluster device failed.
clustering license mismatch with master node	Clustering license does not match that on the master node.
clustering license not available	Clustering license cannot be found.
connection refused by master	Join of a node refused by the master node.
disk in use by another cluster	A disk belongs to a cluster other than the one that a node is joining.
join timed out during reconfiguration	Join of a node has timed out due to reconfiguration taking place in the cluster.
klog update failed	Cannot update kernel log copies during the join of a node.
master aborted during join	Master node aborted while another node was joining the cluster.
protocol version out of range	Cluster protocol version mismatch or unsupported version.
recovery in progress	Volumes that were opened by the node are still recovering.
transition to role failed	Changing the role of a node to be the master failed.
user initiated abort	Node is out of cluster due to an abort initiated by the user or by the cluster monitor.

Table 5-4 Node abort messages (*continued*)

Reason	Description
user initiated stop	Node is out of cluster due to a stop initiated by the user or by the cluster monitor.
vxconfigd is not enabled	The VxVM configuration daemon is not enabled.

See the `vxclustadm(1M)` manual page.

Volume reconfiguration

Volume reconfiguration is the process of creating, changing, and removing VxVM objects such as disk groups, volumes and plexes. In a cluster, all nodes cooperate to perform such operations. The `vxconfigd` daemons play an active role in volume reconfiguration. For reconfiguration to succeed, a `vxconfigd` daemon must be running on each of the nodes.

See [“vxconfigd daemon”](#) on page 162.

A volume reconfiguration transaction is initiated by running a VxVM utility on the master node. The utility contacts the local `vxconfigd` daemon on the master node, which validates the requested change. For example, `vxconfigd` rejects an attempt to create a new disk group with the same name as an existing disk group. The `vxconfigd` daemon on the master node then sends details of the changes to the `vxconfigd` daemons on the slave nodes. The `vxconfigd` daemons on the slave nodes then perform their own checking. For example, each slave node checks that it does not have a private disk group with the same name as the one being created. If the operation involves a new disk, each node checks that it can access that disk. When the `vxconfigd` daemons on all the nodes agree that the proposed change is reasonable, each notifies its kernel. The kernels then cooperate to either commit or to abandon the transaction. Before the transaction can be committed, all of the kernels ensure that no I/O is underway, and block any I/O issued by applications until the reconfiguration is complete. The master node is responsible both for initiating the reconfiguration, and for coordinating the commitment of the transaction. The resulting configuration changes appear to occur simultaneously on all nodes.

If a `vxconfigd` daemon on any node goes away during reconfiguration, all nodes are notified and the operation fails. If any node leaves the cluster, the operation fails unless the master has already committed it. If the master node leaves the cluster, the new master node, which was previously a slave node, completes or fails the operation depending on whether or not it received notification of successful completion from the previous master node. This notification is performed in such a way that if the new master does not receive it, neither does any other slave.

If a node attempts to join a cluster while a volume reconfiguration is being performed, the result of the reconfiguration depends on how far it has progressed. If the kernel has not yet been invoked, the volume reconfiguration is suspended until the node has joined the cluster. If the kernel has been invoked, the node waits until the reconfiguration is complete before joining the cluster.

When an error occurs, such as when a check on a slave fails or a node leaves the cluster, the error is returned to the utility and a message is sent to the console on the master node to identify on which node the error occurred.

vxconfigd daemon

The VxVM configuration daemon, `vxconfigd`, maintains the configuration of VxVM objects. It receives cluster-related instructions from the kernel. A separate copy of `vxconfigd` runs on each node, and these copies communicate with each other over a network. When invoked, a VxVM utility communicates with the `vxconfigd` daemon running on the same node; it does not attempt to connect with `vxconfigd` daemons on other nodes. During cluster startup, the kernel prompts `vxconfigd` to begin cluster operation and indicates whether it is a master node or a slave node.

When a node is initialized for cluster operation, the `vxconfigd` daemon is notified that the node is about to join the cluster and is provided with the following information from the cluster monitor configuration database:

- cluster ID
- node IDs
- master node ID
- role of the node
- network address of the node

On the master node, the `vxconfigd` daemon sets up the shared configuration by importing shared disk groups, and informs the kernel when it is ready for the slave nodes to join the cluster.

On slave nodes, the `vxconfigd` daemon is notified when the slave node can join the cluster. When the slave node joins the cluster, the `vxconfigd` daemon and the VxVM kernel communicate with their counterparts on the master node to set up the shared configuration.

When a node leaves the cluster, the kernel notifies the `vxconfigd` daemon on all the other nodes. The master node then performs any necessary cleanup. If the master node leaves the cluster, the kernels select a new master node and the `vxconfigd` daemons on all nodes are notified of the choice.

The `vxconfigd` daemon also participates in volume reconfiguration.

See [“Volume reconfiguration”](#) on page 161.

vxconfigd daemon recovery

In a cluster, the `vxconfigd` daemons on the slave nodes are always connected to the `vxconfigd` daemon on the master node. If the `vxconfigd` daemon is stopped, volume reconfiguration cannot take place. Other nodes can join the cluster if the `vxconfigd` daemon is not running on the slave nodes.

If the `vxconfigd` daemon stops, different actions are taken depending on which node this occurred:

- If the `vxconfigd` daemon is stopped on the master node, the `vxconfigd` daemons on the slave nodes periodically attempt to rejoin to the master node. Such attempts do not succeed until the `vxconfigd` daemon is restarted on the master. In this case, the `vxconfigd` daemons on the slave nodes have not lost information about the shared configuration, so that any displayed configuration information is correct.
- If the `vxconfigd` daemon is stopped on a slave node, the master node takes no action. When the `vxconfigd` daemon is restarted on the slave, the slave `vxconfigd` daemon attempts to reconnect to the master daemon and to re-acquire the information about the shared configuration. (Neither the kernel view of the shared configuration nor access to shared disks is affected.) Until the `vxconfigd` daemon on the slave node has successfully reconnected to the `vxconfigd` daemon on the master node, it has very little information about the shared configuration and any attempts to display or modify the shared configuration can fail. For example, shared disk groups listed using the `vxldg list` command are marked as `disabled`; when the rejoin completes successfully, they are marked as `enabled`.
- If the `vxconfigd` daemon is stopped on both the master and slave nodes, the slave nodes do not display accurate configuration information until `vxconfigd` is restarted on the master and slave nodes, and the daemons have reconnected.

If the CVM agent for VCS determines that the `vxconfigd` daemon is not running on a node during a cluster reconfiguration, `vxconfigd` is restarted automatically.

Warning: The `-r` reset option to `vxconfigd` restarts the `vxconfigd` daemon and recreates all states from scratch. This option cannot be used to restart `vxconfigd` while a node is joined to a cluster because it causes cluster information to be discarded.

It may sometimes be necessary to restart `vxconfigd` manually in a VCS controlled cluster to resolve a VxVM issue.

To restart vxconfigd manually

- 1 Use the following command to disable failover on any service groups that contain VxVM objects:

```
# hagrps -freeze groupname
```

- 2 Enter the following command to stop and restart the VxVM configuration daemon on the affected node:

```
# vxconfigd -k
```

- 3 Use the following command to re-enable failover for the service groups that you froze in step 1:

```
# hagrps -unfreeze groupname
```

Node shutdown

Although it is possible to shut down the cluster on a node by invoking the shutdown procedure of the node's cluster monitor, this procedure is intended for terminating cluster components after stopping any applications on the node that have access to shared storage. VxVM supports clean node shutdown, which allows a node to leave the cluster gracefully when all access to shared volumes has ceased. The host is still operational, but cluster applications cannot be run on it.

The CVM functionality of VxVM maintains global state information for each volume. This enables VxVM to determine which volumes need to be recovered when a node crashes. When a node leaves the cluster due to a crash or by some other means that is not clean, VxVM determines which volumes may have writes that have not completed and the master node resynchronizes these volumes. It can use dirty region logging (DRL) or FastResync if these are active for any of the volumes.

Clean node shutdown must be used after, or in conjunction with, a procedure to halt all cluster applications. Depending on the characteristics of the clustered application and its shutdown procedure, a successful shutdown can require a lot of time (minutes to hours). For instance, many applications have the concept of draining, where they accept no new work, but complete any work in progress before exiting. This process can take a long time if, for example, a long-running transaction is active.

When the VxVM shutdown procedure is invoked, it checks all volumes in all shared disk groups on the node that is being shut down. The procedure then either continues with the shutdown, or fails for one of the following reasons:

- If all volumes in shared disk groups are closed, VxVM makes them unavailable to applications. Because all nodes are informed that these volumes are closed on the leaving node, no resynchronization is performed.
- If any volume in a shared disk group is open, the shutdown procedure fails. The shutdown procedure can be repeatedly retried until it succeeds. There is no timeout checking in this operation—it is intended as a service that verifies that the clustered applications are no longer active.

Once shutdown succeeds, the node has left the cluster. It is not possible to access the shared volumes until the node joins the cluster again.

Since shutdown can be a lengthy process, other reconfiguration can take place while shutdown is in progress. Normally, the shutdown attempt is suspended until the other reconfiguration completes. However, if it is already too far advanced, the shutdown may complete first.

Cluster shutdown

If all nodes leave a cluster, shared volumes must be recovered when the cluster is next started if the last node did not leave cleanly, or if resynchronization from previous nodes leaving uncleanly is incomplete. CVM automatically handles the recovery and resynchronization tasks when a node joins the cluster.

Dirty region logging in cluster environments

Dirty region logging (DRL) is an optional property of a volume that provides speedy recovery of mirrored volumes after a system failure. DRL is supported in cluster-shareable disk groups. This section provides a brief overview of how DRL behaves in a cluster environment.

In a cluster environment, the VxVM implementation of DRL differs slightly from the normal implementation.

A dirty region log on a system without cluster support has a recovery map and a single active map. A CVM DRL, however, has a single recovery map per cluster and one active map per cluster node.

The dirty region log size in clusters is typically larger than in non-clustered systems, as it must accommodate a recovery map plus active maps for each node in the cluster. The size of each map within the dirty region log is one or more whole blocks.

The `vxassist` command automatically allocates a sufficiently large dirty region log for the size of the volume and the number of nodes.

It is possible to reimport a non-shared disk group (and its volumes) as a shared disk group in a cluster environment. However, the dirty region logs of the imported disk group may be considered invalid and a full recovery may result.

If a shared disk group is imported as a private disk group on a system without cluster support, VxVM considers the logs of the shared volumes to be invalid and conducts a full volume recovery. After the recovery completes, VxVM uses DRL.

The cluster functionality of VxVM can perform a DRL recovery on a non-shared volume. However, if such a volume is moved to a VxVM system with cluster support and imported as shared, the dirty region log is probably too small to accommodate maps for all the cluster nodes. VxVM then marks the log invalid and performs a full recovery anyway. Similarly, moving a DRL volume from a two-node cluster to a four-node cluster can result in too small a log size, which the cluster functionality of VxVM handles with a full volume recovery. In both cases, you must allocate a new log of sufficient size.

See [“Dirty region logging”](#) on page 127.

How DRL works in a cluster environment

When one or more nodes in a cluster crash, DRL must handle the recovery of all volumes that were in use by those nodes when the crashes occurred. On initial cluster startup, all active maps are incorporated into the recovery map during the volume start operation.

Nodes that crash (that is, leave the cluster as dirty) are not allowed to rejoin the cluster until their DRL active maps have been incorporated into the recovery maps on all affected volumes. The recovery utilities compare a crashed node's active maps with the recovery map and make any necessary updates. Only then can the node rejoin the cluster and resume I/O to the volume (which overwrites the active map). During this time, other nodes can continue to perform I/O.

VxVM tracks which nodes have crashed. If multiple node recoveries are underway in a cluster at a given time, VxVM tracks changes in the state of DRL recovery and prevents I/O collisions.

The master node performs volatile tracking of DRL recovery map updates for each volume, and prevents multiple utilities from changing the recovery map simultaneously.

Multiple host failover configurations

Outside the context of CVM, VxVM disk groups can be imported (made available) on only one host at any given time. When a host imports a (private) disk group, the volumes and configuration of that disk group become accessible to the host. If the administrator or system software wants to privately use the same disk group from another host, the host that already has the disk group imported (importing host) must deport (give up access to) the disk group. Once deported, the disk group can be imported by another host.

If two hosts are allowed to access a disk group concurrently without proper synchronization, such as that provided by Oracle RAC, the configuration of the disk group, and possibly the contents of volumes, can be corrupted. Similar corruption can also occur if a file system or database on a raw disk partition is accessed concurrently by two hosts, so this problem is not limited to Veritas Volume Manager.

Import lock

When a host in a non-CVM environment imports a disk group, an import lock is written on all disks in that disk group. The import lock is cleared when the host departs the disk group. The presence of the import lock prevents other hosts from importing the disk group until the importing host has deported the disk group.

Specifically, when a host imports a disk group, the import normally fails if any disks within the disk group appear to be locked by another host. This allows automatic re-importing of disk groups after a reboot (autoimporting) and prevents imports by another host, even while the first host is shut down. If the importing host is shut down without deporting the disk group, the disk group can only be imported by another host by clearing the host ID lock first (discussed later).

The import lock contains a host ID (the host name) reference to identify the importing host and enforce the lock. Problems can therefore arise if two hosts have the same host ID.

Since Veritas Volume Manager uses the host name as the host ID (by default), it is advisable to change the host name of one machine if another machine shares its host name. To change the host name, use the `vxdctl hostid new_hostname` command.

Failover

The import locking scheme works well in an environment where disk groups are not normally shifted from one system to another. However, consider a setup where two hosts, Node A and Node B, can access the drives of a disk group. The disk group is initially imported by Node A, but the administrator wants to access the disk

group from Node B if Node A crashes. Such a failover scenario can be used to provide manual high availability to data, where the failure of one node does not prevent access to data. Failover can be combined with a “high availability” monitor to provide automatic high availability to data: when Node B detects that Node A has crashed or shut down, Node B imports (fails over) the disk group to provide access to the volumes.

Veritas Volume Manager can support failover, but it relies on the administrator or on an external high-availability monitor, such as VCS, to ensure that the first system is shut down or unavailable before the disk group is imported to another system.

See the `vxvg(1M)` manual page.

Corruption of disk group configuration

If `vxvg import` is used with `-C` (clears locks) and/or `-f` (forces import) to import a disk group that is still in use from another host, disk group configuration corruption is likely to occur. Volume content corruption is also likely if a file system or database is started on the imported volumes before the other host crashes or shuts down.

If this kind of corruption occurs, your configuration must typically be rebuilt from scratch and all data be restored from a backup. There are typically numerous configuration copies for each disk group, but corruption nearly always affects all configuration copies, so redundancy does not help in this case.

As long as the configuration backup daemon, `vxconfigbackupd`, is running, VxVM will backup configurations whenever the configuration is changed. By default, backups are stored in `/etc/vx/cbr/bk`. You may also manually backup the configuration using the `vxconfigbackup` utility. The configuration can be rebuilt using the `vxconfigrestore` utility.

See the `vxconfigbackup`, `vxconfigbackupd`, `vxconfigrestore` man pages.

Disk group configuration corruption usually shows up as missing or duplicate records in the configuration databases. This can result in a variety of `vxconfigd` error messages

```
VxVM vxconfigd ERROR
V-5-1-569 Disk group group,Disk disk:
Cannot auto-import group: reason
```

where the *reason* can describe errors such as:

```
Association not resolved
Association count is incorrect
Duplicate record in configuration
Configuration records are inconsistent
```

These errors are typically reported in association with specific disk group configuration copies, but usually apply to all copies. The following is usually displayed along with the error:

```
Disk group has no valid configuration copies
```

If you use the Veritas Cluster Server product, all disk group failover issues can be managed correctly. VCS includes a high availability monitor and includes failover scripts for VxVM, VxFS, and for several popular databases.

The `-t` option to `vxdg` prevents automatic re-imports on reboot and is necessary when used with a host monitor (such as VCS) that controls imports itself, rather than relying on automatic imports by Veritas Volume Manager.

See the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide*.

How Veritas Dynamic Multi-Pathing works

This chapter includes the following topics:

- [How DMP works](#)
- [Veritas Volume Manager co-existence with Oracle Automatic Storage Management \(ASM\) disks](#)

How DMP works

Veritas Dynamic Multi-Pathing (DMP) provides greater availability, reliability, and performance by using path failover and load balancing. This feature is available for multiported disk arrays from various vendors.

Disk arrays can be connected to host systems through multiple paths. To detect the various paths to a disk, DMP uses a mechanism that is specific to each supported array. DMP can also differentiate between different enclosures of a supported array that are connected to the same host system.

See [“Discovering and configuring newly added disk devices”](#) on page 243.

The multi-pathing policy that is used by DMP depends on the characteristics of the disk array.

DMP supports the following standard array types:

Active/Active (A/A)

Allows several paths to be used concurrently for I/O. Such arrays allow DMP to provide greater I/O throughput by balancing the I/O load uniformly across the multiple paths to the LUNs. In the event that one path fails, DMP automatically routes I/O over the other available paths.

Asymmetric Active/Active (A/A-A)	A/A-A or Asymmetric Active/Active arrays can be accessed through secondary storage paths with little performance degradation. The behavior is similar to ALUA, except that it does not support those SCSI commands which an ALUA array supports.
Asymmetric Logical Unit Access (ALUA)	DMP supports all variants of ALUA.
Active/Passive (A/P)	Allows access to its LUNs (logical units; real disks or virtual disks created using hardware) via the primary (active) path on a single controller (also known as an access port or a storage processor) during normal operation. In implicit failover mode (or autotrespass mode), an A/P array automatically fails over by scheduling I/O to the secondary (passive) path on a separate controller if the primary path fails. This passive port is not used for I/O until the active port fails. In A/P arrays, path failover can occur for a single LUN if I/O fails on the primary path. This policy supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.
Active/Passive in explicit failover mode or non-autotrespass mode (A/P-F)	The appropriate command must be issued to the array to make the LUNs fail over to the secondary path. This policy supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.

Active/Passive with LUN group failover (A/P-G)	For Active/Passive arrays with LUN group failover (A/P-G arrays), a group of LUNs that are connected through a controller is treated as a single failover entity. Unlike A/P arrays, failover occurs at the controller level, and not for individual LUNs. The primary controller and the secondary controller are each connected to a separate group of LUNs. If a single LUN in the primary controller's LUN group fails, all LUNs in that group fail over to the secondary controller. This policy supports concurrent I/O and load balancing by having multiple primary paths into a controller. This functionality is provided by a controller with multiple ports, or by the insertion of a SAN switch between an array and a controller. Failover to the secondary (passive) path occurs only if all the active primary paths fail.


An array policy module (APM) may define array types to DMP in addition to the standard types for the arrays that it supports.

Veritas Storage Foundation Cluster File System High Availability uses DMP metanodes (DMP nodes) to access disk devices connected to the system. For each disk in a supported array, DMP maps one node to the set of paths that are connected to the disk. Additionally, DMP associates the appropriate multi-pathing policy for the disk array with the node.

For disks in an unsupported array, DMP maps a separate node to each path that is connected to a disk. The raw and block devices for the nodes are created in the directories `/dev/vx/rdmp` and `/dev/vx/dmp` respectively.

[Figure 6-1](#) shows how DMP sets up a node for a disk in a supported disk array.


Figure 6-1 How DMP represents multiple physical paths to a disk as one node



DMP implements a disk device naming scheme that allows you to recognize to which array a disk belongs.

Figure 6-2 shows an example where two paths, `c1t99d0` and `c2t99d0`, exist to a single disk in the enclosure, but VxVM uses the single DMP node, `enc0_0`, to access it.

Figure 6-2 Example of multi-pathing for a disk enclosure in a SAN environment



See [“About enclosure-based naming”](#) on page 174.

See [“Changing the disk device naming scheme”](#) on page 327.

Device discovery

Device discovery is the term used to describe the process of discovering the disks that are attached to a host. This feature is important for DMP because it needs to support a growing number of disk arrays from a number of vendors. In conjunction with the ability to discover the devices attached to a host, the Device Discovery service enables you to add support for new disk arrays. The Device discovery uses a facility called the Device Discovery Layer (DDL).

The DDL enables you to add support for new disk arrays without the need for a reboot.

This means that you can dynamically add a new disk array to a host, and run a command which scans the operating system's device tree for all the attached disk devices, and reconfigures DMP with the new device database.

See [“How to administer the Device Discovery Layer”](#) on page 249.


About enclosure-based naming

Enclosure-based naming provides an alternative to operating system-based device naming. In a Storage Area Network (SAN) that uses Fibre Channel switches, information about disk location provided by the operating system may not correctly indicate the physical location of the disks. Enclosure-based naming allows SFCFSHA to access enclosures as separate physical entities. By configuring redundant copies of your data on separate enclosures, you can safeguard against failure of one or more enclosures.

Enclosure-based naming allows disk devices to be named for enclosures rather than for the controllers through which they are accessed. For example, `c##t##d##s##` naming assigns controller-based device names to disks in separate enclosures that are connected to the same host controller.

[Figure 6-3](#) shows a typical SAN environment where host controllers are connected to multiple enclosures through a Fibre Channel switch.

Figure 6-3 Example configuration for disk enclosures connected via a fibre channel switch



In such a configuration, enclosure-based naming can be used to refer to each disk within an enclosure. For example, the device names for the disks in enclosure `enc0` are named `enc0_0`, `enc0_1`, and so on. The main benefit of this scheme is that it allows you to quickly determine where a disk is physically located in a large SAN configuration.

In most disk arrays, you can use hardware-based storage management to represent several physical disks as one LUN to the operating system. In such cases, VxVM also sees a single logical disk device rather than its component disks. For this reason, when reference is made to a disk within an enclosure, this disk may be either a physical disk or a LUN.


Another important benefit of enclosure-based naming is that it enables VxVM to avoid placing redundant copies of data in the same enclosure. This is a good thing to avoid as each enclosure can be considered to be a separate fault domain. For example, if a mirrored volume were configured only on the disks in enclosure `enc1`, the failure of the cable between the switch and the enclosure would make the entire volume unavailable.

If required, you can replace the default name that SFCFSA assigns to an enclosure with one that is more meaningful to your configuration.

See [“Renaming an enclosure”](#) on page 300.

Figure 6-4 shows a High Availability (HA) configuration where redundant-loop access to storage is implemented by connecting independent controllers on the host to separate switches with independent paths to the enclosures.

Figure 6-4 Example HA configuration using multiple switches to provide redundant loop access



Such a configuration protects against the failure of one of the host controllers (`c1` and `c2`), or of the cable between the host and one of the switches. In this example, each disk is known by the same name to VxVM for all of the paths over which it can be accessed.

For example, the disk device `enc0_0` represents a single disk for which two different paths are known to the operating system, such as `c1t99d0` and `c2t99d0`.

See [“Changing the disk device naming scheme”](#) on page 327.

To take account of fault domains when configuring data redundancy, you can control how mirrored volumes are laid out across enclosures.

How DMP monitors I/O on paths

In VxVM prior to release 5.0, DMP had one kernel daemon (`errorld`) that performed error processing, and another (`restored`) that performed path restoration activities.

From release 5.0, DMP maintains a pool of kernel threads that are used to perform such tasks as error processing, path restoration, statistics collection, and SCSI

request callbacks. The `vxddmpadm gettune` command can be used to provide information about the threads. The name `restored` has been retained for backward compatibility.

One kernel thread responds to I/O failures on a path by initiating a probe of the host bus adapter (HBA) that corresponds to the path. Another thread then takes the appropriate action according to the response from the HBA. The action taken can be to retry the I/O request on the path, or to fail the path and reschedule the I/O on an alternate path.

The restore kernel task is woken periodically (typically every 5 minutes) to check the health of the paths, and to resume I/O on paths that have been restored. As some paths may suffer from intermittent failure, I/O is only resumed on a path if the path has remained healthy for a given period of time (by default, 5 minutes). DMP can be configured with different policies for checking the paths.

See [“Configuring DMP path restoration policies”](#) on page 305.

The statistics-gathering task records the start and end time of each I/O request, and the number of I/O failures and retries on each path. DMP can be configured to use this information to prevent the SCSI driver being flooded by I/O requests. This feature is known as I/O throttling.

If an I/O request relates to a mirrored volume, VxVM specifies the FAILFAST flag.

In such cases, DMP does not retry failed I/O requests on the path, and instead marks the disks on that path as having failed.

See [“Path failover mechanism”](#) on page 177.

See [“I/O throttling”](#) on page 178.

Path failover mechanism

DMP enhances system availability when used with disk arrays having multiple paths. In the event of the loss of a path to a disk array, DMP automatically selects the next available path for I/O requests without intervention from the administrator.

DMP is also informed when a connection is repaired or restored, and when you add or remove devices after the system has been fully booted (provided that the operating system recognizes the devices correctly).

If required, the response of DMP to I/O failure on a path can be tuned for the paths to individual arrays. DMP can be configured to time out an I/O request either after a given period of time has elapsed without the request succeeding, or after a given number of retries on a path have failed.

See [“Configuring DMP's response to I/O failures”](#) on page 301.

Subpaths Failover Group (SFG)

A subpaths failover group (SFG) represents a group of paths which could fail and restore together. When an I/O error is encountered on a path in an SFG, DMP does proactive path probing on the other paths of that SFG as well. This behavior adds greatly to the performance of path failover thus improving IO performance. Currently the criteria followed by DMP to form the subpaths failover groups is to bundle the paths with the same endpoints from the host to the array into one logical storage failover group.

See [“Configuring DMP's Subpaths Failover Groups \(SFG\)”](#) on page 304.

Low Impact Path Probing (LIPP)

The restore daemon in DMP keeps probing the LUN paths periodically. This behavior helps DMP to keep the path states up-to-date even though IO activity is not there on the paths. Low Impact Path Probing adds logic to the restore daemon to optimize the number of the probes performed while the path status is being updated by the restore daemon. This optimization is achieved with the help of the logical subpaths failover groups. With LIPP logic in place, DMP probes only limited number of paths within an SFG, instead of probing all the paths in an SFG. Based on these probe results, DMP determines the states of all the paths in that SFG.

See [“Configuring DMP's Low Impact Path Probing”](#) on page 303.

I/O throttling

If I/O throttling is enabled, and the number of outstanding I/O requests builds up on a path that has become less responsive, DMP can be configured to prevent new I/O requests being sent on the path either when the number of outstanding I/O requests has reached a given value, or a given time has elapsed since the last successful I/O request on the path. While throttling is applied to a path, the new I/O requests on that path are scheduled on other available paths. The throttling is removed from the path if the HBA reports no error on the path, or if an outstanding I/O request on the path succeeds.

See [“Configuring the DMP I/O throttling mechanism”](#) on page 302.

Load balancing

By default, Veritas Dynamic Multi-Pathing (DMP) uses the Minimum Queue I/O policy for load balancing across paths for Active/Active (A/A), Active/Passive (A/P), Active/Passive with explicit failover (A/P-F) and Active/Passive with group failover (A/P-G) disk arrays. Load balancing maximizes I/O throughput by using the total bandwidth of all available paths. I/O is sent down the path which has the minimum outstanding I/Os.

For A/P disk arrays, I/O is sent down the primary paths. If all of the primary paths fail, I/O is switched over to the available secondary paths. As the continuous transfer of ownership of LUNs from one controller to another results in severe I/O slowdown, load balancing across primary and secondary paths is not performed for A/P disk arrays unless they support concurrent I/O.

For A/P, A/P-F and A/P-G arrays, load balancing is performed across all the currently active paths as is done for A/A arrays.

You can change the I/O policy for the paths to an enclosure or disk array.

See [“Specifying the DMP I/O policy”](#) on page 292.

Disabling Sun Multipath IO (MPxIO)

Plan for system downtime for the following procedure.

This procedure describes steps to migrate devices from MPxIO control to DMP.

The migration steps involve system downtime on a host due to the following:

- Need to stop applications
- Need to stop the VCS services if using VCS
- The procedure involves one or more host reboots

To take devices out of MPxIO control and enable DMP on the devices

- 1 Stop the applications that use MPxIO devices.
- 2 Unmount all the file systems that use MPxIO devices.
- 3 Disable MPxIO using the following command.

```
# stmsboot -d
```

- 4 Reboot the system.
- 5 Mount the file systems.
- 6 Restart the applications.

Dynamic Reconfiguration

Dynamic Reconfiguration (DR) is a feature that is available on some high-end enterprise systems. It allows some components (such as CPUs, memory, and other controllers or I/O boards) to be reconfigured while the system is still running. The reconfigured component might be handling the disks controlled by VxVM.

See [“About enabling and disabling I/O for controllers and storage processors”](#) on page 264.

See [“About online dynamic reconfiguration”](#) on page 310.

About booting from DMP devices

When the root disk is placed under VxVM control, it is automatically accessed as a DMP device with one path if it is a single disk, or with multiple paths if the disk is part of a multiported disk array. By encapsulating and mirroring the root disk, system reliability is enhanced against loss of one or more of the existing physical paths to a disk.

The `vxrootadm` utility can be used to clone the encapsulated boot disk group to a self-contained disk group specified by the user.

Note: The SAN bootable LUN must be controlled by DMP. PowerPath and MPxIO control of SAN bootable LUNs is not supported.

DMP in a clustered environment

Note: You need an additional license to use the cluster feature of VxVM. Clustering is only supported for VxVM.

In a clustered environment where Active/Passive type disk arrays are shared by multiple hosts, all nodes in the cluster must access the disk via the same physical storage controller port. Accessing a disk via multiple paths simultaneously can severely degrade I/O performance (sometimes referred to as the ping-pong effect). Path failover on a single cluster node is also coordinated across the cluster so that all the nodes continue to share the same physical path.

Prior to release 4.1 of VxVM, the clustering and DMP features could not handle automatic failback in A/P arrays when a path was restored, and did not support failback for explicit failover mode arrays. Failback could only be implemented manually by running the `vxctl enable` command on each cluster node after the path failure had been corrected. From release 4.1, failback is now an automatic cluster-wide operation that is coordinated by the master node. Automatic failback in explicit failover mode arrays is also handled by issuing the appropriate low-level command.

Note: Support for automatic failback of an A/P array requires that an appropriate ASL (and APM, if required) is installed on the system.

See [“Discovering disks and dynamically adding disk arrays”](#) on page 245.

For Active/Active type disk arrays, any disk can be simultaneously accessed through all available physical paths to it. In a clustered environment, the nodes do not need to access a disk via the same physical path.

See [“How to administer the Device Discovery Layer”](#) on page 249.

See [“Configuring array policy modules”](#) on page 307.

About enabling or disabling controllers with shared disk groups

See [“How to administer the Device Discovery Layer”](#) on page 249.

See [“DMP in a clustered environment”](#) on page 180.

Prior to release 5.0, VxVM did not allow enabling or disabling of paths or controllers connected to a disk that is part of a shared Veritas Volume Manager disk group. From VxVM 5.0 onward, such operations are supported on shared DMP nodes in a cluster.

Veritas Volume Manager co-existence with Oracle Automatic Storage Management (ASM) disks

ASM disks are the disks used by Oracle Automatic Storage Management software. Veritas Volume Manager (VxVM) co-exists with Oracle ASM disks, by recognizing the disks as the type Oracle ASM. VxVM protects ASM disks from any operations that may overwrite the disk. VxVM classifies and displays the ASM disks as ASM format disks. You cannot initialize an ASM disk, or perform any VxVM operations that may overwrite the disk.

If the disk is claimed as an ASM disk, disk initialization commands fail with an appropriate failure message. The `vxdisk init` command and the `vxdisksetup` command fail, even if the `force` option is specified. The `vxprivutil` command also fails for disks under ASM control, to prevent any on-disk modification of the ASM device.

If the target disk is under ASM control, any rootability operations that overwrite the target disk fail. A message indicates that the disk is already in use as an ASM disk. The rootability operations include operations to create a VM root image (`vxcp_lvmroot` command), create a VM root mirror (`vxrootmir` command), or restore the LVM root image (`vxres_lvmroot` command). The `vxdestroy_lvmroot` command also fails for ASM disks, since the target disk is not under LVM control as expected.

Disks that ASM accessed previously but that no longer belong to an ASM disk group are called FORMER ASM disks. If you remove an ASM disk from ASM control, VxVM labels the disk as a FORMER ASM disk. VxVM enforces the same restrictions

for FORMER ASM disks as for ASM disks, to enable ASM to reuse the disk in the future. To use a FORMER ASM disk with VxVM, you must clean the disk of ASM information after you remove the disk from ASM control. If a disk initialization command is issued on a FORMER ASM disk, the command fails. A message indicates that the disk must be cleaned up before the disk can be initialized for use with VxVM.

To remove a FORMER ASM disk from ASM control for use with VxVM

- 1 Clean the disk with the `dd` command to remove all ASM identification information on it. For example:

```
dd if=/dev/zero of=/dev/rdisk/<wholedisk|partition> count=1 bs=1024
```

Where *wholedisk* is a disk name in the format: `cxydz`

Where *partition* is a partition name in the format: `cxydzsn`

- 2 Perform a disk scan:

```
# vxdisk scandisks
```

To view the ASM disks

- ◆ You can use either of the following commands to display ASM disks:

The `vxdisk list` command displays the disk type as ASM.

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
Disk_0s2	auto:LVM	-	-	LVM
Disk_1	auto:ASM	-	-	ASM
EVA4K6K0_0	auto	-	-	online
EVA4K6K0_1	auto	-	-	online

The `vxdisk classify` command classifies and displays ASM disks as Oracle ASM.

```
# vxdisk -d classify disk=c1t0d5
```

```
device:      c1t0d5
status:      CLASSIFIED
type:        Oracle ASM
groupid:     -
hostname:    -
domainid:    -
centralhost: -
```

Specify the `-f` option to the `vxdisk classify` command, to perform a full scan of the OS devices.

To check if a particular disk is under ASM control

- ◆ Use the `vxisasm` utility to check if a particular disk is under ASM control.

```
# /etc/vx/bin/vxisasm 3pardata0_2799
```

```
3pardata0_2799  ACTIVE
```

```
# /etc/vx/bin/vxisasm 3pardata0_2798
```

```
3pardata0_2798  FORMER
```

Alternatively, use the `vxisforeign` utility to check if the disk is under control of any foreign software like LVM or ASM:

```
# /etc/vx/bin/vxisforeign 3pardata0_2799
```

```
3pardata0_2799  ASM    ACTIVE
```

```
# /etc/vx/bin/vxisforeign 3pardata0_2798
```

```
3pardata0_2798  ASM    FORMER
```

Provisioning storage

- [Chapter 7. Provisioning new storage](#)
- [Chapter 8. Advanced allocation methods for configuring storage](#)
- [Chapter 9. Creating and mounting VxFS file systems](#)
- [Chapter 10. Extent attributes](#)

Provisioning new storage

This chapter includes the following topics:

- [Provisioning new storage](#)
- [Growing the existing storage by adding a new LUN](#)
- [Growing the existing storage by growing the LUN](#)
- [Displaying SFCFSA information with vxlist](#)

Provisioning new storage

The following procedure describes how to provision new storage. If you are provisioning Storage Foundation on thin storage, you should understand how Storage Foundation works with thin storage.

See [“About thin optimization solutions in Storage Foundation Cluster File System High Availability”](#) on page 674.

To provision new storage

- 1 Set up the LUN. See the documentation for your storage array for information about how to create, mask, and bind the LUN.
- 2 Initialize the LUNs for Veritas Volume Manager (VxVM), using one of the following commands.

The recommended method is to use the `vxdisksetup` command.

```
# vxdisksetup -i 3PARDATA0_1
# vxdisk init 3PARDATA0_1
```

- 3 Add the LUN to a disk group.
 - If you do not have a disk group for your LUN, create the disk group:

```
# vxdg init dg1 dev1=3PARDATA0_1
```

- If you already have a disk group for your LUN, add the LUN to the disk group:

```
# vxdg -g dg1 adddisk 3PARDATA0_1
```

- 4 Create the volume on the LUN:

```
# vxassist -b -g dg1 make vol1 100g 3PARDATA0_1
```

- 5 Create a Veritas File System (VxFS) file system on the volume:

```
# mkfs -F vxfs /dev/vx/rdisk/dg1/vol1
```

- 6 Create a mount point on the file system:

```
# mkdir mount1
```

- 7 Mount the file system:

```
# mount -F vxfs /dev/vx/dsk/dg1/vol1 /mount1
```

Growing the existing storage by adding a new LUN

The following procedure describes how to grow the existing storage by adding a new LUN.

To grow the existing storage by adding a new LUN

- 1 Create and set up the LUN.
- 2 Add the LUN to the disk group.

```
# vxdg -g dg1 adddisk 3PARDATA0_2
```

- 3 Grow the volume and the file system to the desired size. For example:

```
# vxresize -b -F vxfs -g dg1 vol1 100g
```

Growing the existing storage by growing the LUN

The following procedure describes how to grow the existing storage by growing a LUN.

To grow the existing storage by growing a LUN

- 1 Grow the existing LUN. See the documentation for your storage array for information about how to create, mask, and bind the LUN.
- 2 Make VxVM aware of the new LUN size.

```
# vxdisk -g dg1 resize c0t1d0s4
```

See [“Dynamic LUN expansion”](#) on page 318.

- 3 Calculate the new maximum volume size:

```
# vxassist -g dg1 -b maxgrow vol1
```

- 4 Grow the volume and the file system to the desired size:

```
# vxresize -b -F vxfs -g dg1 vol1 100g
```

Displaying SFCFSHA information with vxlist

The `vxlist` command is a new display command that provides a consolidated view of the SFCFSHA configuration. The `vxlist` command consolidates information from Veritas Volume Manager (VxVM) and Veritas File System (VxFS). The `vxlist` command provides various options to display information. For example, use the following form of the command to display file system information including information about the volume, disk group, and so on. In previous releases, you needed to run at least two commands to retrieve the following information.

```
# /opt/VRTSsfmh/bin/vxlist fs
```

TY	FS	FSTYPE	SIZE	FREE	%USED	DEVICE_PATH	MOUNT_POINT
fs	/	ext3	65.20g	51.70g	17%	/dev/sda1	/
fs	mnt	vxfs	19.84g	9.96g	49%	/dev/vx/dsk/bardg/vol1	/mnt

For help on the `vxlist` command, enter the following command:

```
# vxlist -H
```

See the `vxlist(1m)` manual page.

Advanced allocation methods for configuring storage

This chapter includes the following topics:

- [Customizing allocation behavior](#)
- [Creating volumes of a specific layout](#)
- [Creating a volume on specific disks](#)
- [Creating volumes on specific media types](#)
- [Specifying ordered allocation of storage to volumes](#)
- [Site-based allocation](#)
- [Changing the read policy for mirrored volumes](#)

Customizing allocation behavior

By default, the `vxassist` command creates volumes on any available storage that meets basic requirements. The `vxassist` command seeks out available disk space and allocates it in the configuration that conforms to the layout specifications and that offers the best use of free space. The `vxassist` command creates the required plexes and subdisks using only the basic attributes of the desired volume as input.

If you are provisioning Storage Foundation on thin storage, you should understand how Storage Foundation works with thin storage.

See [“About thin optimization solutions in Storage Foundation Cluster File System High Availability”](#) on page 674.

Additionally, when you modify existing volumes using the `vxassist` command, the `vxassist` command automatically modifies underlying or associated objects. The `vxassist` command uses default values for many volume attributes, unless you provide specific values to the command line. You can customize the default behavior of the `vxassist` command by customizing the default values.

See [“Setting default values for vxassist”](#) on page 189.

The `vxassist` command creates volumes in a default disk group according to the default rules. To use a different disk group, specify the `-g diskgroup` option to the `vxassist` command.

See [“Rules for determining the default disk group”](#) on page 191.

If you want to assign particular characteristics for a certain volume, you can specify additional attributes on the `vxassist` command line. Examples are stripe unit width, number of columns in a RAID-5 or stripe volume, number of mirrors, number of logs, and log type.

For details of available `vxassist` keywords and attributes, refer to the `vxassist(1M)` manual page.

You can use allocation attributes to specify the following types of allocation behavior:

Layouts for the volumes	See “Creating volumes of a specific layout” on page 198.
Media types	See “Creating volumes on specific media types” on page 206.
Specific disks, subdisks, plexes locations	See “Creating a volume on specific disks” on page 204.
Ordered allocation	See “Specifying ordered allocation of storage to volumes” on page 206.
Site-based allocation	See “Site-based allocation” on page 209.
Setting the read policy	See “Changing the read policy for mirrored volumes” on page 209.

Setting default values for vxassist

The default values that the `vxassist` command uses may be specified in the file `/etc/default/vxassist`. The defaults listed in this file take effect if you do not override them on the command line, or in an alternate defaults file that you specify

using the `-d` option. A default value specified on the command line always takes precedence. `vxassist` also has a set of built-in defaults that it uses if it cannot find a value defined elsewhere.

You must create the `/etc/default` directory and the `vxassist` default file if these do not already exist on your system.

The format of entries in a defaults file is a list of attribute-value pairs separated by new lines. These attribute-value pairs are the same as those specified as options on the `vxassist` command line.

See the `vxassist(1M)` manual page.

To display the default attributes held in the file `/etc/default/vxassist`, use the following form of the `vxassist` command:

```
# vxassist help showattrs
```

The following is a sample `vxassist` defaults file:

```
# By default:
# create unmirrored, unstriped volumes
# allow allocations to span drives
# with RAID-5 create a log, with mirroring don't create a log
# align allocations on cylinder boundaries
    layout=nomirror,nostripe,span,nocontig,raid5log,noregionlog,
    diskalign

# use the fsgen usage type, except when creating RAID-5 volumes
    usetype=fsgen

# allow only root access to a volume
    mode=u=rw,g=,o=
    user=root
    group=root

# when mirroring, create two mirrors
    nmirror=2

# for regular striping, by default create between 2 and 8 stripe
# columns
    max_nstripe=8
    min_nstripe=2

# for RAID-5, by default create between 3 and 8 stripe columns
    max_nraid5stripe=8
    min_nraid5stripe=3
```

```
# by default, create 1 log copy for both mirroring and RAID-5 volumes
    nregionlog=1
    nraid5log=1

# by default, limit mirroring log lengths to 32Kbytes
    max_regionloglen=32k

# use 64K as the default stripe unit size for regular volumes
    stripe_stwid=64k

# use 16K as the default stripe unit size for RAID-5 volumes
    raid5_stwid=16k
```

Rules for determining the default disk group

You should use the `-g` option to specify a disk group to Veritas Volume Manager (VxVM) commands that accept this option. If you do not specify the disk group, VxVM applies rules in the following order until it determines a disk group name:

- Use the default disk group name that is specified by the environment variable `VXVM_DEFAULTDG`. This variable can also be set to one of the reserved system-wide disk group names: `bootdg`, `defaultdg`, or `nodg`.
See [“Displaying the system-wide boot disk group”](#) on page 191.
If the variable is undefined, the following rule is applied.
- Use the disk group that has been assigned to the system-wide default disk group alias, `defaultdg`.
See [“Displaying and specifying the system-wide default disk group”](#) on page 192.
If this alias is undefined, the following rule is applied.
- If the operation can be performed without requiring a disk group name (for example, an edit operation on disk access records), do so.

If none of these rules succeeds, the requested operation fails.

Warning: In releases of VxVM prior to 4.0, a subset of commands tried to determine the disk group by searching for the object name that was being operated upon by a command. This functionality is no longer supported. Scripts that rely on determining the disk group from an object name may fail.

Displaying the system-wide boot disk group

To display the currently defined system-wide boot disk group, use the following command:

```
# vxdg bootdg
```

See the `vxdg(1M)` manual page.

Displaying and specifying the system-wide default disk group

Veritas Volume Manager (VxVM) enables you to define a system-wide default disk group.

To display the currently defined system-wide default disk group, use the following command:

```
# vxdg defaultdg
```

If a default disk group has not been defined, `nodg` is displayed. You can also use the following command to display the default disk group:

```
# vxprint -Gng defaultdg 2>/dev/null
```

In this case, if there is no default disk group, nothing is displayed.

Use the following command to specify the name of the disk group that is aliased by `defaultdg`:

```
# vxdctl defaultdg diskgroup
```

If `bootdg` is specified as the argument to this command, the default disk group is set to be the same as the currently defined system-wide boot disk group.

If `nodg` is specified as the argument to the `vxdctl defaultdg` command, the default disk group is undefined.

The specified disk group is not required to exist on the system.

See the `vxdctl(1M)` manual page.

See the `vxdg(1M)` manual page.

Using rules and persistent attributes to make volume allocation more efficient

The `vxassist` command lets you create a set of volume allocation rules and define it with a single name. When you specify this name in your volume allocation request, all the attributes that are defined in this rule are honored when `vxassist` creates the volume.

Creating volume allocation rules has the following benefits:

- Rules streamline your typing and reduce errors. You can define relatively complex allocation rules once in a single location and reuse them.

- Rules let you standardize behaviors in your environment, including across a set of servers.

For example, you can create allocation rules so that a set of servers can standardize their storage tiering. Suppose you had the following requirements:

Tier 1	Enclosure mirroring between a specific set of array types
Tier 2	Non-mirrored striping between a specific set of array types
Tier 0	Select solid-state drive (SSD) storage

You can create rules for each volume allocation requirement and name the rules tier1, tier2, and tier0.

You can also define rules so that each time you create a volume for a particular purpose, it's created with the same attributes. For example, to create the volume for a production database, you can create a rule called productiondb. To create standardized volumes for home directories, you can create a rule called homedir. To standardize your high performance index volumes, you can create a rule called dbindex.

Understanding persistent attributes

The `vxassist` command lets you record certain volume allocation attributes for a volume. These attributes are called persistent attributes. You can record the attributes which would be useful in later allocation operations on the volume. Useful attributes include volume grow and enclosure mirroring. You can also restrict allocation to storage that has a particular property (such as the enclosure type, disk tag, or media type). On the other hand, volume length is not useful, and generally neither is a specific list of disks.

The persistent attributes can be retrieved and applied to the allocation requests (with possible modifications) for the following operations:

- volume grow or shrink
- move
- relayout
- mirror
- add a log

Persistent attributes let you record carefully described allocation attributes at the time of volume creation and retain them for future allocation operations on the volume. Also, you can modify, enhance, or even discard the persistent attributes. For example, you can add and retain a separation rule for a volume that is originally

not mirrored. Alternatively, you can temporarily suspend a volume allocation which has proven too restrictive or discard it to allow a needed allocation to succeed.

Rule file format

When you create rules, you do not define them in the `/etc/default/vxassist` file. You create the rules in another file and add the path information to `/etc/default/vxassist`. By default, a rule file is loaded from `/etc/default/vxsf_rules`. You can override this location in `/etc/default/vxassist` with the attribute `rulefile=/path/rule_file_name`. You can also specify additional rule files on the command line.

A rule file uses the following conventions:

- Blank lines are ignored.
- Use the pound sign, `#`, to begin a comment.
- Use C language style quoting for the strings that may include embedded spaces, new lines, or tabs. For example, use quotes around the text for the `description` attribute.
- Separate tokens with a space.
- Use braces for a rule that is longer than one line.

Within the rule file, a volume allocation rule has the following format:

```
volume rule rulename vxassist_attributes
```

This syntax defines a rule named *rulename* which is a short-hand for the listed `vxassist` attributes. Rules can reference other rules using an attribute of `rule=rulename[,rulename,...]`, which adds all the attributes from that rule into the rule currently being defined. The attributes you specify in a rule definition override any conflicting attributes that are in a rule that you specify by reference. You can add a description to a rule with the attribute `description=description_text`.

The following is a basic rule file. The first rule in the file, `base`, defines the `logtype` and `persist` attributes. The remaining rules in the file – `tier0`, `tier1`, and `tier2` – reference this rule and also define their own tier-specific attributes. Referencing a rule lets you define attributes in one place and reuse them in other rules.

```
# Create tier 1 volumes mirrored between disk arrays, tier 0 on SSD,  
# and tier 2 as unmirrored. Always use FMR DCO objects.  
volume rule base { logtype=dco persist=yes }  
volume rule tier0 { rule=base mediatype:ssd tier=tier0 }  
volume rule tier1 { rule=base mirror=enclosure tier=tier1 }  
volume rule tier2 { rule=base tier=tier2 }
```

The following rule file contains a more complex definition which runs across several lines.

```
volume rule appXdb_storage {
    description="Create storage for the database of Application X"
    rule=base
    siteconsistent=yes
    mirror=enclosure
}
```

By default, a rule file is loaded from `/etc/default/vxsf_rules`. You can override this location in `/etc/default/vxassist`. You can also specify additional rule files on the command line.

Using rules to create a volume

When you use the `vxassist` command to create a volume, you can include the rule name on the command line. For example, the content of the `vxsf_rules` file is as follows:

```
volume rule basic { logtype=dco }
volume rule tier1 {
    rule=basic
    layout=mirror
    tier=tier1
}
```

In the following example, when you create the volume `vol1` in disk group `dg3`, you can specify the `tier1` rule on the command line. In addition to the attributes you enter on the command line, `vol1` is given the attributes that you defined in `tier1`.

```
vxassist -g dg3 make vol1 200m rule=tier1
```

The following `vxprint` command displays the attributes of disk group `dg3`. The output includes the new volume, `vol1`.

```
vxprint -g dg3
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	dg3	dg3	-	-	-	-	-	-
dm	ibm_ds8x000_0266	ibm_ds8x000_0266	-	2027264	-	-	-	-
dm	ibm_ds8x000_0267	ibm_ds8x000_0267	-	2027264	-	-	-	-
dm	ibm_ds8x000_0268	ibm_ds8x000_0268	-	2027264	-	-	-	-
v	vol1	fsgen	ENABLED	409600	-	ACTIVE	-	-

```

pl vol1-01      vol1          ENABLED  409600  -        ACTIVE  -        -
sd ibm_ds8x000_0266-01 vol1-01 ENABLED  409600  0        -        -        -
pl vol1-02      vol1          ENABLED  409600  -        ACTIVE  -        -
sd ibm_ds8x000_0267-01 vol1-02 ENABLED  409600  0        -        -        -
dc vol1_dco     vol1          -        -        -        -        -        -
v  vol1_dcl     gen           ENABLED  144      -        ACTIVE  -        -
pl vol1_dcl-01  vol1_dcl     ENABLED  144      -        ACTIVE  -        -
sd ibm_ds8x000_0266-02 vol1_dcl-01 ENABLED  144  0        -        -        -
pl vol1_dcl-02  vol1_dcl     ENABLED  144      -        ACTIVE  -        -
sd ibm_ds8x000_0267-02 vol1_dcl-02 ENABLED  144  0        -        -        -

```

The following `vxassist` command confirms that `vol1` is in `tier1`. The application of rule `tier1` was successful.

```

vxassist -g dg3 listtag
TY  NAME                DISKGROUP          TAG
=====
v   vol1                dg3                vxfs.placement_class.tier1

```

Using persistent attributes

You can define volume allocation attributes so they can be reused in subsequent operations. These attributes are called persistent attributes, and they are stored in a set of hidden volume tags. The `persist` attribute determines whether an attribute persists, and how the current command might use or modify preexisting persisted attributes. You can specify persistence rules in defaults files, in rules, or on the command line. For more information, see the `vxassist` manual page.

To illustrate how persistent attributes work, we'll use the following `vxsf_rules` file. It contains a rule, `rule1`, which defines the `mediatype` attribute. This rule also uses the `persist` attribute to make the `mediatype` attribute persistent.

```

# cat /etc/default/vxsf_rules
volume rule rule1 { mediatype:ssd persist=extended }

```

The following command confirms that LUNs `ibm_ds8x000_0266` and `ibm_ds8x000_0268` are solid-state disk (SSD) devices.

```

# vxdisk listtag
DEVICE                NAME                VALUE
ibm_ds8x000_0266      vxmediatype        ssd
ibm_ds8x000_0268      vxmediatype        ssd

```

The following command creates a volume, `vol1`, in the disk group `dg3`. `rule1` is specified on the command line, so those attributes are also applied to `vol1`.


```
# vxassist -g dg3 make vol1 100m rule=rule1
```

The following command shows that the volume `vol1` is created off the SSD device `ibm_ds8x000_0266` as specified in rule1.

```
# vxprint -g dg3
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	dg3	dg3	-	-	-	-	-	-
dm	ibm_ds8x000_0266	ibm_ds8x000_0266	-	2027264	-	-	-	-
dm	ibm_ds8x000_0267	ibm_ds8x000_0267	-	2027264	-	-	-	-
dm	ibm_ds8x000_0268	ibm_ds8x000_0268	-	2027264	-	-	-	-
v	vol1	fsgen	ENABLED	204800	-	ACTIVE	-	-
pl	vol1-01	vol1	ENABLED	204800	-	ACTIVE	-	-
sd	ibm_ds8x000_0266-01	vol1-01	ENABLED	204800	0	-	-	-

The following command displays the attributes that are defined in rule1.

```
# vxassist -g dg3 help showattr rule=rule1
```

```
alloc=mediatype:ssd
```

```
persist=extended
```

If no persistent attributes are defined, the following command grows `vol1` on hard disk drive (HDD) devices. However, at the beginning of this section, `mediatype:ssd` was defined as a persistent attribute. Therefore, the following command honors this original intent and grows the volume on SSD devices.

```
# vxassist -g dg3 growby vol1 1g
```

The following `vxprint` command confirms that the volume was grown on SSD devices.

```
# vxprint -g dg3
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	dg3	dg3	-	-	-	-	-	-
dm	ibm_ds8x000_0266	ibm_ds8x000_0266	-	2027264	-	-	-	-
dm	ibm_ds8x000_0267	ibm_ds8x000_0267	-	2027264	-	-	-	-
dm	ibm_ds8x000_0268	ibm_ds8x000_0268	-	2027264	-	-	-	-
v	vol1	fsgen	ENABLED	2301952	-	ACTIVE	-	-
pl	vol1-01	vol1	ENABLED	2301952	-	ACTIVE	-	-
sd	ibm_ds8x000_0266-01	vol1-01	ENABLED	2027264	0	-	-	-
sd	ibm_ds8x000_0268-01	vol1-01	ENABLED	274688	2027264	-	-	-

Creating volumes of a specific layout

VxVM enables you to create volumes of various layouts. You can specify an attribute to indicate the type of layout you want to create. The following sections include details for each of the following types:

- * Creating a mirrored volume
- * Creating a striped volume
- * Creating a RAID-5 volume

Types of volume layouts

VxVM allows you to create volumes with the following layout types:

Concatenated	A volume whose subdisks are arranged both sequentially and contiguously within a plex. Concatenation allows a volume to be created from multiple regions of one or more disks if there is not enough space for an entire volume on a single region of a disk. If a single LUN or disk is split into multiple subdisks, and each subdisk belongs to a unique volume, this is called carving. See “Concatenation, spanning, and carving” on page 106.
Striped	A volume with data spread evenly across multiple disks. Stripes are equal-sized fragments that are allocated alternately and evenly to the subdisks of a single plex. There must be at least two subdisks in a striped plex, each of which must exist on a different disk. Throughput increases with the number of disks across which a plex is striped. Striping helps to balance I/O load in cases where high traffic areas exist on certain subdisks. See “Striping (RAID-0)” on page 108.
Mirrored	A volume with multiple data plexes that duplicate the information contained in a volume. Although a volume can have a single data plex, at least two are required for true mirroring to provide redundancy of data. For the redundancy to be useful, each of these data plexes should contain disk space from different disks. See “Mirroring (RAID-1)” on page 111.

RAID-5	A volume that uses striping to spread data and parity evenly across multiple disks in an array. Each stripe contains a parity stripe unit and data stripe units. Parity can be used to reconstruct data if one of the disks fails. In comparison to the performance of striped volumes, write throughput of RAID-5 volumes decreases since parity information needs to be updated each time data is modified. However, in comparison to mirroring, the use of parity to implement data redundancy reduces the amount of space required. See “RAID-5 (striping with parity)” on page 114.
Mirrored-stripe	A volume that is configured as a striped plex and another plex that mirrors the striped one. This requires at least two disks for striping and one or more other disks for mirroring (depending on whether the plex is simple or striped). The advantages of this layout are increased performance by spreading data across multiple disks and redundancy of data. See “Striping plus mirroring (mirrored-stripe or RAID-0+1)” on page 112.
Layered Volume	A volume constructed from other volumes. Non-layered volumes are constructed by mapping their subdisks to VM disks. Layered volumes are constructed by mapping their subdisks to underlying volumes (known as storage volumes), and allow the creation of more complex forms of logical layout. See “Layered volumes” on page 119. The following layouts are examples of layered volumes: ■ Striped-mirror A striped-mirror volume is created by configuring several mirrored volumes as the columns of a striped volume. This layout offers the same benefits as a non-layered mirrored-stripe volume. In addition, it provides faster recovery as the failure of single disk does not force an entire striped plex offline. See “Mirroring plus striping (striped-mirror, RAID-1+0 or RAID-10)” on page 113. ■ Concatenated-mirror A concatenated-mirror volume is created by concatenating several mirrored volumes. This provides faster recovery as the failure of a single disk does not force the entire mirror offline.

Creating a mirrored volume

A mirrored volume provides data redundancy by containing more than one copy of its data. Each copy (or mirror) is stored on different disks from the original copy of

the volume and from other mirrors. Mirroring a volume ensures that its data is not lost if a disk in one of its component mirrors fails.

A mirrored volume requires space to be available on at least as many disks in the disk group as the number of mirrors in the volume.

If you specify `layout=mirror`, `vxassist` determines the best layout for the mirrored volume. Because the advantages of the layouts are related to the size of the volume, `vxassist` selects the layout based on the size of the volume. For smaller volumes, `vxassist` uses the simpler mirrored concatenated (`mirror-concat`) layout. For larger volumes, `vxassist` uses the more complex concatenated mirror (`concat-mirror`) layout. The attribute `stripe-mirror-col-split-trigger-pt` controls the selection. Volumes that are smaller than `stripe-mirror-col-split-trigger-pt` are created as `mirror-concat`, and volumes that are larger are created as `concat-mirror`. By default, the attribute `stripe-mirror-col-split-trigger-pt` is set to one gigabyte. The value can be set in `/etc/default/vxassist`. If there is a reason to implement a particular layout, you can specify `layout=mirror-concat` or `layout=concat-mirror` to implement the desired layout.

To create a new mirrored volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \  
    layout=mirror [nmirror=number] [init=active]
```

Specify the `-b` option if you want to make the volume immediately available for use.

For example, to create the mirrored volume, `volmir`, in the disk group, `mydg`, use the following command:

```
# vxassist -b -g mydg make volmir 5g layout=mirror
```

To create a volume with 3 instead of the default of 2 mirrors, modify the command to read:

```
# vxassist -b -g mydg make volmir 5g layout=mirror nmirror=3
```

Creating a mirrored-concatenated volume

A mirrored-concatenated volume mirrors several concatenated plexes. To create a concatenated-mirror volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \  
    layout=mirror-concat [nmirror=number]
```

Specify the `-b` option if you want to make the volume immediately available for use.

Alternatively, first create a concatenated volume, and then mirror it.

See [“Adding a mirror to a volume”](#) on page 838.

Creating a concatenated-mirror volume

A concatenated-mirror volume is an example of a layered volume which concatenates several underlying mirror volumes. To create a concatenated-mirror volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \  
    layout=concat-mirror [nmirror=number]
```

Specify the `-b` option if you want to make the volume immediately available for use.

Creating a striped volume

A striped volume contains at least one plex that consists of two or more subdisks located on two or more physical disks. A striped volume requires space to be available on at least as many disks in the disk group as the number of columns in the volume.

See [“Striping \(RAID-0\)”](#) on page 108.

To create a striped volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length layout=stripe
```

Specify the `-b` option if you want to make the volume immediately available for use.

For example, to create the 10-gigabyte striped volume `volzebra`, in the disk group, `mydg`, use the following command:

```
# vxassist -b -g mydg make volzebra 10g layout=stripe
```

This creates a striped volume with the default stripe unit size (64 kilobytes) and the default number of stripes (2).

You can specify the disks on which the volumes are to be created by including the disk names on the command line. For example, to create a 30-gigabyte striped volume on three specific disks, `mydg03`, `mydg04`, and `mydg05`, use the following command:

```
# vxassist -b -g mydg make stripevol 30g layout=stripe \  
    mydg03 mydg04 mydg05
```

To change the number of columns or the stripe width, use the `ncolumn` and `stripeunit` modifiers with `vxassist`. For example, the following command creates a striped volume with 5 columns and a 32-kilobyte stripe size:

```
# vxassist -b -g mydg make stripevol 30g layout=stripe \  
    stripeunit=32k ncol=5
```

Creating a mirrored-stripe volume

A mirrored-stripe volume mirrors several striped data plexes. A mirrored-stripe volume requires space to be available on at least as many disks in the disk group as the number of mirrors multiplied by the number of columns in the volume.

To create a striped-mirror volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \  
  layout=mirror-stripe [nmirror=number_of_mirrors] \  
  [ncol=number_of_columns] [stripewidth=size]
```

Specify the `-b` option if you want to make the volume immediately available for use.

Alternatively, first create a striped volume, and then mirror it. In this case, the additional data plexes may be either striped or concatenated.

See [“Adding a mirror to a volume”](#) on page 838.

Creating a striped-mirror volume

A striped-mirror volume is an example of a layered volume which stripes several underlying mirror volumes. A striped-mirror volume requires space to be available on at least as many disks in the disk group as the number of columns multiplied by the number of stripes in the volume.

To create a striped-mirror volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \  
  layout=stripe-mirror [nmirror=number_of_mirrors] \  
  [ncol=number_of_columns] [stripewidth=size]
```

Specify the `-b` option if you want to make the volume immediately available for use.

By default, VxVM attempts to create the underlying volumes by mirroring subdisks rather than columns if the size of each column is greater than the value for the attribute `stripe-mirror-col-split-trigger-pt` that is defined in the `vxassist` defaults file.

If there are multiple subdisks per column, you can choose to mirror each subdisk individually instead of each column. To mirror at the subdisk level, specify the layout as `stripe-mirror-sd` rather than `stripe-mirror`. To mirror at the column level, specify the layout as `stripe-mirror-col` rather than `stripe-mirror`.

Creating a RAID-5 volume

A RAID-5 volume requires space to be available on at least as many disks in the disk group as the number of columns in the volume. Additional disks may be required for any RAID-5 logs that are created.

Note: VxVM supports the creation of RAID-5 volumes in private disk groups, but not in shareable disk groups in a cluster environment.

You can create RAID-5 volumes by using either the `vxassist` command (recommended) or the `vxmake` command. Both approaches are described below.

A RAID-5 volume contains a RAID-5 data plex that consists of three or more subdisks located on three or more physical disks. Only one RAID-5 data plex can exist per volume. A RAID-5 volume can also contain one or more RAID-5 log plexes, which are used to log information about data and parity being written to the volume.

See [“RAID-5 \(striping with parity\)”](#) on page 114.

Warning: Do not create a RAID-5 volume with more than 8 columns because the volume will be unrecoverable in the event of the failure of more than one disk.

To create a RAID-5 volume, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length layout=raid5 \  
    [ncol=number_of_columns] [stripewidth=size] [nlog=number] \  
    [loglen=log_length]
```

Specify the `-b` option if you want to make the volume immediately available for use.

For example, to create the RAID-5 volume `volraid` together with 2 RAID-5 logs in the disk group, `mydg`, use the following command:

```
# vxassist -b -g mydg make volraid 10g layout=raid5 nlog=2
```

This creates a RAID-5 volume with the default stripe unit size on the default number of disks. It also creates two RAID-5 logs rather than the default of one log.

If you require RAID-5 logs, you must use the `logdisk` attribute to specify the disks to be used for the log plexes.

RAID-5 logs can be concatenated or striped plexes, and each RAID-5 log associated with a RAID-5 volume has a complete copy of the logging information for the volume. To support concurrent access to the RAID-5 array, the log should be several times the stripe size of the RAID-5 plex.

It is suggested that you configure a minimum of two RAID-5 log plexes for each RAID-5 volume. These log plexes should be located on different disks. Having two RAID-5 log plexes for each RAID-5 volume protects against the loss of logging information due to the failure of a single disk.

If you use ordered allocation when creating a RAID-5 volume on specified storage, you must use the `logdisk` attribute to specify on which disks the RAID-5 log plexes should be created. Use the following form of the `vxassist` command to specify the disks from which space for the logs is to be allocated:

```
# vxassist [-b] [-g diskgroup] -o ordered make volumelength \
  layout=raid5 [ncol=number_columns] [nlog=number] \
  [loglen=log_length] logdisk=disk[,disk,...] \
  storage_attributes
```

For example, the following command creates a 3-column RAID-5 volume with the default stripe unit size on disks `mydg04`, `mydg05` and `mydg06`. It also creates two RAID-5 logs on disks `mydg07` and `mydg08`.

```
# vxassist -b -g mydg -o ordered make volraid 10g layout=raid5 \
  ncol=3 nlog=2 logdisk=mydg07,mydg08 mydg04 mydg05 mydg06
```

The number of logs must equal the number of disks that is specified to `logdisk`.

See [“Specifying ordered allocation of storage to volumes”](#) on page 206.

See the `vxassist(1M)` manual page.

It is possible to add more logs to a RAID-5 volume at a later time.

Creating a volume on specific disks

VxVM automatically selects the disks on which each volume resides, unless you specify otherwise. If you want a volume to be created on specific disks, you must designate those disks to VxVM. More than one disk can be specified.

To create a volume on a specific disk or disks, use the following command:

```
# vxassist [-b] [-g diskgroup] make volume length \
  [layout=layout] diskname ...
```

Specify the `-b` option if you want to make the volume immediately available for use.

For example, to create the volume `volspec` with length 5 gigabytes on disks `mydg03` and `mydg04`, use the following command:

```
# vxassist -b -g mydg make volspec 5g mydg03 mydg04
```


The `vxassist` command allows you to specify storage attributes. These give you control over the devices, including disks, controllers and targets, which `vxassist` uses to configure a volume. For example, you can specifically exclude disk `mydg05`.

Note: The `!` character is a special character in some shells. The following examples show how to escape it in a bash shell.

```
# vxassist -b -g mydg make volspec 5g \!mydg05
```

The following example excludes all disks that are on controller `c2`:

```
# vxassist -b -g mydg make volspec 5g \!ctlr:c2
```

This example includes only disks on controller `c1` except for target `t5`:

```
# vxassist -b -g mydg make volspec 5g ctlr:c1 \!target:c1t5
```

If you want a volume to be created using only disks from a specific disk group, use the `-g` option to `vxassist`, for example:

```
# vxassist -g bigone -b make volmega 20g bigone10 bigone11
```

or alternatively, use the `diskgroup` attribute:

```
# vxassist -b make volmega 20g diskgroup=bigone bigone10 \
    bigone11
```

Any storage attributes that you specify for use must belong to the disk group. Otherwise, `vxassist` will not use them to create a volume.

You can also use storage attributes to control how `vxassist` uses available storage, for example, when calculating the maximum size of a volume, when growing a volume or when removing mirrors or logs from a volume. The following example excludes disks `dgrp07` and `dgrp08` when calculating the maximum size of RAID-5 volume that `vxassist` can create using the disks in the disk group `dg`:

```
# vxassist -b -g dgrp maxsize layout=raid5 nlog=2 \!dgrp07 \!dgrp08
```

It is also possible to control how volumes are laid out on the specified storage.

See [“Specifying ordered allocation of storage to volumes”](#) on page 206.

If you are using VxVM in conjunction with Veritas SANPoint Control 2.0, you can specify how `vxassist` should use the available storage groups when creating volumes.

See [“Configuration of volumes on SAN storage”](#) on page 142.

See the `vxassist(1M)` manual page.

`vxassist` also lets you select disks based on disk tags. The following command only includes disks that have a `tier1` disktag.

```
# vxassist -g dg3 make vol3 1g disktag:tier1
```

Creating volumes on specific media types

When you create a volume, you can specify the media type for the volume. The supported media types are Hard Disk Drives (HDD) or Solid State Devices (SSD). The SSD media type requires disk group 150 or greater. The default is HDD.

To specify a media type, specify the `vxassist` command with the `mediatype` attribute. If no `mediatype` is specified, the volume allocates storage only from the HDD devices.

Specifying ordered allocation of storage to volumes

Ordered allocation gives you complete control of space allocation. It requires that the number of disks that you specify to the `vxassist` command must match the number of disks that are required to create a volume. The order in which you specify the disks to `vxassist` is also significant.

If you specify the `-o ordered` option to `vxassist` when creating a volume, any storage that you also specify is allocated in the following order:


- Concatenate disks.
- Form columns.
- Form mirrors.

For example, the following command creates a mirrored-stripe volume with 3 columns and 2 mirrors on 6 disks in the disk group, `mydg`:

```
# vxassist -b -g mydg -o ordered make mirstrvol 10g \  
  layout=mirror-stripe ncol=3 mydg01 mydg02 mydg03 \  
  mydg04 mydg05 mydg06
```

This command places columns 1, 2 and 3 of the first mirror on disks `mydg01`, `mydg02` and `mydg03` respectively, and columns 1, 2 and 3 of the second mirror on disks `mydg04`, `mydg05` and `mydg06` respectively.

[Figure 8-1](#) shows an example of using ordered allocation to create a mirrored-stripe volume.


Figure 8-1 Example of using ordered allocation to create a mirrored-stripe volume

For layered volumes, `vxassist` applies the same rules to allocate storage as for non-layered volumes. For example, the following command creates a striped-mirror volume with 2 columns:

```
# vxassist -b -g mydg -o ordered make strmirvol 10g \
  layout=stripe-mirror ncol=2 mydg01 mydg02 mydg03 mydg04
```

This command mirrors column 1 across disks `mydg01` and `mydg03`, and column 2 across disks `mydg02` and `mydg04`.

Figure 8-2 shows an example of using ordered allocation to create a striped-mirror volume.

Figure 8-2 Example of using ordered allocation to create a striped-mirror volume


Additionally, you can use the `col_switch` attribute to specify how to concatenate space on the disks into columns. For example, the following command creates a mirrored-stripe volume with 2 columns:

```
# vxassist -b -g mydg -o ordered make strmir2vol 10g \
  layout=mirror-stripe ncol=2 col_switch=3g,2g \
  mydg01 mydg02 mydg03 mydg04 mydg05 mydg06 mydg07 mydg08
```

This command allocates 3 gigabytes from `mydg01` and 2 gigabytes from `mydg02` to column 1, and 3 gigabytes from `mydg03` and 2 gigabytes from `mydg04` to column 2. The mirrors of these columns are then similarly formed from disks `mydg05` through `mydg08`.

Figure 8-3 shows an example of using concatenated disk space to create a mirrored-stripe volume.

Figure 8-3 Example of using concatenated disk space to create a mirrored-stripe volume




Other storage specification classes for controllers, enclosures, targets and trays can be used with ordered allocation. For example, the following command creates a 3-column mirrored-stripe volume between specified controllers:

```
# vxassist -b -g mydg -o ordered make mirstr2vol 80g \
  layout=mirror-stripe ncol=3 \
  ctrlr:c1 ctrlr:c2 ctrlr:c3 ctrlr:c4 ctrlr:c5 ctrlr:c6
```

This command allocates space for column 1 from disks on controllers `c1`, for column 2 from disks on controller `c2`, and so on.

Figure 8-4 shows an example of using storage allocation to create a mirrored-stripe volume across controllers.

Figure 8-4 Example of storage allocation used to create a mirrored-stripe volume across controllers



There are other ways in which you can control how `vxassist` lays out mirrored volumes across controllers.

Site-based allocation

In a Remote Mirror configuration (also known as a campus cluster or stretch cluster), the hosts and storage of a cluster are divided between two or more sites. These sites are typically connected via a redundant high-capacity network that provides access to storage and private link communication between the cluster nodes.

Configure the disk group in a Remote Mirror site to be site-consistent. When you create volumes in such a disk group, the volumes are mirrored across all sites by default.

See [“About sites and remote mirrors”](#) on page 464.

Changing the read policy for mirrored volumes

VxVM offers the choice of the following read policies on the data plexes in a mirrored volume:

`round`

Reads each plex in turn in “round-robin” fashion for each nonsequential I/O detected. Sequential access causes only one plex to be accessed. This approach takes advantage of the drive or controller read-ahead caching policies.

<code>prefer</code>	Reads first from a plex that has been named as the preferred plex.
<code>select</code>	Chooses a default policy based on plex associations to the volume. If the volume has an enabled striped plex, the <code>select</code> option defaults to preferring that plex; otherwise, it defaults to round-robin. For disk group versions 150 or higher and if there is a SSD based plex available, it will be preferred over other plexes.
<code>siteread</code>	Reads preferentially from plexes at the locally defined site. This method is the default policy for volumes in disk groups where site consistency has been enabled. For disk group versions 150 or higher and if the local site has a SSD based plex, it will be preferred. See “About site consistency” on page 468.
<code>split</code>	Divides the read requests and distributes them across all the available plexes.

Note: You cannot set the read policy on a RAID-5 volume.

To set the read policy to `round`, use the following command:

```
# vxvol [-g diskgroup] rdpol round volume
```

For example, to set the read policy for the volume `vol101` in disk group `mydg` to round-robin, use the following command:

```
# vxvol -g mydg rdpol round vol101
```

To set the read policy to `prefer`, use the following command:

```
# vxvol [-g diskgroup] rdpol prefer volume preferred_plex
```

For example, to set the policy for `vol101` to read preferentially from the plex `vol101-02`, use the following command:

```
# vxvol -g mydg rdpol prefer vol101 vol101-02
```

To set the read policy to `select`, use the following command:

```
# vxvol [-g diskgroup] rdpol select volume
```

Creating and mounting VxFS file systems

This chapter includes the following topics:

- [Creating a VxFS file system](#)
- [Converting a file system to VxFS](#)
- [Mounting a VxFS file system](#)
- [Unmounting a file system](#)
- [Resizing a file system](#)
- [Displaying information on mounted file systems](#)
- [Identifying file system types](#)
- [Monitoring free space](#)

Creating a VxFS file system

The `mkfs` command creates a VxFS file system by writing to a special character device file. The special character device is a location or character device node of a particular storage device. `mkfs` builds a file system with a root directory and a `lost+found` directory.

Before running `mkfs`, you must create the target device. Refer to your operating system documentation for more information. If you are using a logical device (such as a VxVM volume), see the VxVM documentation for instructions on device initialization.

See the `mkfs(1M)` and `mkfs_vxfs(1M)` manual pages.

When you create a file system with the `mkfs` command, you can select the following characteristics:

- [Block size](#)
- [Intent log size](#)

To create a file system

- ◆ Use the `mkfs` command to create a file system:

```
mkfs [-F vxfs] [-m] [generic_options] [-o specific_options] \
special [size]
```

<code>-F vxfs</code>	Specifies the VxFS file system type.
<code>-m</code>	Displays the command line that was used to create the file system. The file system must already exist. This option enables you to determine the parameters used to construct the file system.
<i>generic_options</i>	Options common to most other file system types.
<code>-o specific_options</code>	Options specific to VxFS.
<code>-o N</code>	Displays the geometry of the file system and does not write to the device.
<code>-o largefiles</code>	Allows users to create files larger than two gigabytes. The default option is <code>largefiles</code> .
<i>special</i>	Specifies the special device file location or character device node of a particular storage device.
<i>size</i>	Specifies the number of 512-byte sectors in the file system. If <i>size</i> is not specified, <code>mkfs</code> determines the size of the special device.

Block size

The unit of allocation in VxFS is an extent. Unlike some other UNIX file systems, VxFS does not make use of block fragments for allocation because storage is allocated in extents that consist of one or more blocks.

You specify the block size when creating a file system by using the `mkfs -o bsize` option. The block size cannot be altered after the file system is created. The smallest available block size for VxFS is 1K. The default block size is 1024 bytes for file systems smaller than 1 TB, and 8192 bytes for file systems 1 TB or larger.

Choose a block size based on the type of application being run. For example, if there are many small files, a 1K block size may save space. For large file systems, with relatively few files, a larger block size is more appropriate. Larger block sizes use less disk space in file system overhead, but consume more space for files that are not a multiple of the block size. The easiest way to judge which block sizes provide the greatest system efficiency is to try representative system loads against various sizes and pick the fastest. For most applications, it is best to use the default values.

For 64-bit kernels, the block size and disk layout version determine the maximum size of the file system you can create.

See [“About disk layouts”](#) on page 939.

Intent log size

You specify the intent log size when creating a file system by using the `mkfs -o logsize` option. You can dynamically increase or decrease the intent log size using the `logsize` option of the `fsadm` command. The `mkfs` utility uses a default intent log size of 64 megabytes. The default size is sufficient for most workloads. If the system is used as an NFS server or for intensive synchronous write workloads, performance may be improved using a larger log size.

With larger intent log sizes, recovery time is proportionately longer and the file system may consume more system resources (such as memory) during normal operation.

There are several system performance benchmark suites for which VxFS performs better with larger log sizes. As with block sizes, the best way to pick the log size is to try representative system loads against various sizes and pick the fastest.

Example of creating a file system

The following example creates a VxFS file system of 12288 sectors in size on a VxVM volume.

To create a VxFS file system

1 Create the file system:

```
# mkfs -F vxfs /dev/vx/rdisk/diskgroup/volume 12288
version 9 layout
12288 sectors, 6144 blocks of size 1024, log size 512 blocks
largefiles supported
```

2 Mount the newly created file system:

```
# mount -F vxfs /dev/vx/rdisk/diskgroup/volume /mnt1
```

Converting a file system to VxFS

The `vxfsconvert` command can be used to convert a UFS file system to a VxFS file system.

See the `vxfsconvert(1M)` manual page.

To convert a UFS file system to a VxFS file system

- ◆ Use the `vxfsconvert` command to convert a UFS file system to VxFS:

```
vxfsconvert [-l logsize] [-s size] [-efnNvyY] special
```

<code>-e</code>	Estimates the amount of space required to complete the conversion.
<code>-f</code>	Displays the list of supported file system types.
<code>-l logsize</code>	Specifies the size of the file system intent log.
<code>-n N</code>	Assumes a no response to all questions asked by <code>vxfsconvert</code> .
<code>-s size</code>	Directs <code>vxfsconvert</code> to use free disk space past the current end of the file system to store VxFS metadata.
<code>-v</code>	Specifies verbose mode.
<code>-y Y</code>	Assumes a yes response to all questions asked by <code>vxfsconvert</code> .
<i>special</i>	Specifies the name of the character (raw) device that contains the file system to convert.

Example of converting a file system

The following example converts a UFS file system to a VxFS file system with an intent log size of 16384 blocks.

To convert a UFS file system to a VxFS file system

- ◆ Convert the file system:

```
# vxfsconvert -l 16384 /dev/vx/rdisk/diskgroup/volume
```

Mounting a VxFS file system

In addition to the standard mount mode (`delaylog` mode), Veritas File System (VxFS) provides the following mount options for you to specify other modes of operation:

- [log mount option](#)
- [delaylog mount option](#)
- [tmplog mount option](#)
- [logiosize mount option](#)
- [nodatainlog mount option](#)
- [blkclear mount option](#)
- [mincache mount option](#)
- [convosync mount option](#)
- [ioerror mount option](#)
- [largefiles|nolargefiles mount option](#)
- [cio mount option](#)
- [mntlock mount option](#)
- [ckptautomnt mount option](#)

Caching behavior can be altered with the `mincache` option, and the behavior of `O_SYNC` and `D_SYNC` writes can be altered with the `convosync` option.

See the `fcntl(2)` manual page.

The `delaylog` and `tmplog` modes can significantly improve performance. The improvement over `log` mode is typically about 15 to 20 percent with `delaylog`; with `tmplog`, the improvement is even higher. Performance improvement varies, depending on the operations being performed and the workload. Read/write intensive

loads should show less improvement, while file system structure intensive loads, such as `mkdir`, `create`, and `rename`, may show over 100 percent improvement. The best way to select a mode is to test representative system loads against the logging modes and compare the performance results.

Most of the modes can be used in combination. For example, a desktop machine might use both the `blkclear` and `mincache=closesync` modes.

The `mount` command automatically runs the VxFS `fsck` command to clean up the intent log if the `mount` command detects a dirty log in the file system. This functionality is only supported on file systems mounted on a Veritas Volume Manager (VxVM) volume.

See the `mount_vxfs(1M)` manual page.

log mount option

In log mode, all system calls other than `write(2)`, `writew(2)`, and `pwrite(2)` are guaranteed to be persistent after the system call returns to the application.

The `rename(2)` system call flushes the source file to disk to guarantee the persistence of the file data before renaming it. In both the `log` and `delaylog` modes, the `rename` is also guaranteed to be persistent when the system call returns. This benefits shell scripts and programs that try to update a file atomically by writing the new file contents to a temporary file and then renaming it on top of the target file.

delaylog mount option

The default logging mode is `delaylog`. In `delaylog` mode, the effects of most system calls other than `write(2)`, `writew(2)`, and `pwrite(2)` are guaranteed to be persistent approximately 3 seconds after the system call returns to the application. Contrast this with the behavior of most other file systems in which most system calls are not persistent until approximately 30 seconds or more after the call has returned. Fast file system recovery works with this mode.

The `rename(2)` system call flushes the source file to disk to guarantee the persistence of the file data before renaming it. In the `log` and `delaylog` modes, the `rename` is also guaranteed to be persistent when the system call returns. This benefits shell scripts and programs that try to update a file atomically by writing the new file contents to a temporary file and then renaming it on top of the target file.

tmplog mount option

In `tmplog` mode, the effects of system calls have persistence guarantees that are similar to those in `delaylog` mode. In addition, enhanced flushing of delayed

extending writes is disabled, which results in better performance but increases the chances of data being lost or uninitialized data appearing in a file that was being actively written at the time of a system failure. This mode is only recommended for temporary file systems. Fast file system recovery works with this mode.

Note: The term "effects of system calls" refers to changes to file system data and metadata caused by the system call, excluding changes to `st_atime`.

See the `stat(2)` manual page.

Logging mode persistence guarantees

In all logging modes, VxFS is fully POSIX compliant. The effects of the `fsync(2)` and `fdatasync(2)` system calls are guaranteed to be persistent after the calls return. The persistence guarantees for data or metadata modified by `write(2)`, `writew(2)`, or `pwrite(2)` are not affected by the logging mount options. The effects of these system calls are guaranteed to be persistent only if the `O_SYNC`, `O_DSYNC`, `VX_DSYNC`, or `VX_DIRECT` flag, as modified by the `convosync=` mount option, has been specified for the file descriptor.

The behavior of NFS servers on a VxFS file system is unaffected by the `log` and `tmplog` mount options, but not `delaylog`. In all cases except for `tmplog`, VxFS complies with the persistency requirements of the NFS v2 and NFS v3 standard. Unless a UNIX application has been developed specifically for the VxFS file system in `log` mode, it expects the persistence guarantees offered by most other file systems and experiences improved robustness when used with a VxFS file system mounted in `delaylog` mode. Applications that expect better persistence guarantees than that offered by most other file systems can benefit from the `log`, `mincache=`, and `closesync` mount options. However, most commercially available applications work well with the default VxFS mount options, including the `delaylog` mode.

logiosize mount option

The `logiosize=size` option enhances the performance of storage devices that employ a read-modify-write feature. If you specify `logiosize` when you mount a file system, VxFS writes the intent log in the least `size` bytes or a multiple of `size` bytes to obtain the maximum performance from such devices.

See the `mount_vxfs(1M)` manual page.

The values for `size` can be 512, 1024, 2048, 4096, or 8192.

nodatainlog mount option

Use the `nodatainlog` mode on systems with disks that do not support bad block revectoring. Usually, a VxFS file system uses the intent log for synchronous writes. The inode update and the data are both logged in the transaction, so a synchronous write only requires one disk write instead of two. When the synchronous write returns to the application, the file system has told the application that the data is already written. If a disk error causes the metadata update to fail, then the file must be marked bad and the entire file is lost.

If a disk supports bad block revectoring, then a failure on the data update is unlikely, so logging synchronous writes should be allowed. If the disk does not support bad block revectoring, then a failure is more likely, so the `nodatainlog` mode should be used.

A `nodatainlog` mode file system is approximately 50 percent slower than a standard mode VxFS file system for synchronous writes. Other operations are not affected.

blkclear mount option

The `blkclear` mode is used in increased data security environments. The `blkclear` mode guarantees that uninitialized storage never appears in files. The increased integrity is provided by clearing extents on disk when they are allocated within a file. This mode does not affect extending writes. A `blkclear` mode file system is approximately 10 percent slower than a standard mode VxFS file system, depending on the workload.

mincache mount option

The mincache mode has the following suboptions:

- `mincache=closesync`
- `mincache=direct`
- `mincache=dsync`
- `mincache=unbuffered`
- `mincache=tmpcache`

The `mincache=closesync` mode is useful in desktop environments where users are likely to shut off the power on the machine without halting it first. In this mode, any changes to the file are flushed to disk when the file is closed.

To improve performance, most file systems do not synchronously update data and inode changes to disk. If the system crashes, files that have been updated within the past minute are in danger of losing data. With the `mincache=closesync` mode,

if the system crashes or is switched off, only open files can lose data. A `mincache=closesync` mode file system could be approximately 15 percent slower than a standard mode VxFS file system, depending on the workload.

The following describes where to use the mincache modes:

- The `mincache=direct`, `mincache=unbuffered`, and `mincache=dsync` modes are used in environments where applications have reliability problems caused by the kernel buffering of I/O and delayed flushing of non-synchronous I/O.
- The `mincache=direct` and `mincache=unbuffered` modes guarantee that all non-synchronous I/O requests to files are handled as if the `VX_DIRECT` or `VX_UNBUFFERED` caching advisories had been specified.
- The `mincache=dsync` mode guarantees that all non-synchronous I/O requests to files are handled as if the `VX_DSYNC` caching advisory had been specified. Refer to the `vxfsio(7)` manual page for explanations of `VX_DIRECT`, `VX_UNBUFFERED`, and `VX_DSYNC`, as well as for the requirements for direct I/O.
- The `mincache=direct`, `mincache=unbuffered`, and `mincache=dsync` modes also flush file data on close as `mincache=closesync` does.

Because the `mincache=direct`, `mincache=unbuffered`, and `mincache=dsync` modes change non-synchronous I/O to synchronous I/O, throughput can substantially degrade for small to medium size files with most applications. Since the `VX_DIRECT` and `VX_UNBUFFERED` advisories do not allow any caching of data, applications that normally benefit from caching for reads usually experience less degradation with the `mincache=dsync` mode. `mincache=direct` and `mincache=unbuffered` require significantly less CPU time than buffered I/O.

If performance is more important than data integrity, you can use the `mincache=tmpcache` mode. The `mincache=tmpcache` mode disables special delayed extending write handling, trading off less integrity for better performance. Unlike the other `mincache` modes, `tmpcache` does not flush the file to disk the file is closed. When the `mincache=tmpcache` option is used, bad data can appear in a file that was being extended when a crash occurred.

convosync mount option

The `convosync` (convert osync) mode has the following suboptions:

- `convosync=closesync`

Note: The `convosync=closesync` mode converts synchronous and data synchronous writes to non-synchronous writes and flushes the changes to the file to disk when the file is closed.

- `convosync=delay`
- `convosync=direct`
- `convosync=dsync`

Note: The `convosync=dsync` option violates POSIX guarantees for synchronous I/O.

- `convosync=unbuffered`

The `convosync=delay` mode causes synchronous and data synchronous writes to be delayed rather than to take effect immediately. No special action is performed when closing a file. This option effectively cancels any data integrity guarantees normally provided by opening a file with `O_SYNC`.

See the `open(2)`, `fcntl(2)`, and `vxfsio(7)` manual pages.

Warning: Be very careful when using the `convosync=closesync` or `convosync=delay` mode because they actually change synchronous I/O into non-synchronous I/O. Applications that use synchronous I/O for data reliability may fail if the system crashes and synchronously written data is lost.

The `convosync=dsync` mode converts synchronous writes to data synchronous writes.

As with `closesync`, the `direct`, `unbuffered`, and `dsync` modes flush changes to the file to disk when it is closed. These modes can be used to speed up applications that use synchronous I/O. Many applications that are concerned with data integrity specify the `O_SYNC` `fcntl` in order to write the file data synchronously. However, this has the undesirable side effect of updating inode times and therefore slowing down performance. The `convosync=dsync`, `convosync=unbuffered`, and `convosync=direct` modes alleviate this problem by allowing applications to take advantage of synchronous writes without modifying inode times as well.

Before using `convosync=dsync`, `convosync=unbuffered`, or `convosync=direct`, make sure that all applications that use the file system do not require synchronous inode time updates for `O_SYNC` writes.

ioerror mount option

This mode sets the policy for handling I/O errors on a mounted file system. I/O errors can occur while reading or writing file data or metadata. The file system can respond to these I/O errors either by halting or by gradually degrading. The `ioerror`

option provides five policies that determine how the file system responds to the various errors. All policies limit data corruption, either by stopping the file system or by marking a corrupted inode as bad.

The policies are as follows:

- [disable policy](#)
- [nodisable policy](#)
- [wdisable policy and mwdisable policy](#)
- [mdisable policy](#)

disable policy

If `disable` is selected, VxFS disables the file system after detecting any I/O error. You must then unmount the file system and correct the condition causing the I/O error. After the problem is repaired, run `fsck` and mount the file system again. In most cases, replay `fsck` to repair the file system. A full `fsck` is required only in cases of structural damage to the file system's metadata. Select `disable` in environments where the underlying storage is redundant, such as RAID-5 or mirrored disks.

nodisable policy

If `nodisable` is selected, when VxFS detects an I/O error, it sets the appropriate error flags to contain the error, but continues running. Note that the degraded condition indicates possible data or metadata corruption, not the overall performance of the file system.

For file data read and write errors, VxFS sets the `VX_DATAIOERR` flag in the super-block. For metadata read errors, VxFS sets the `VX_FULLFSCK` flag in the super-block. For metadata write errors, VxFS sets the `VX_FULLFSCK` and `VX_METAIOERR` flags in the super-block and may mark associated metadata as bad on disk. VxFS then prints the appropriate error messages to the console.

You should stop the file system as soon as possible and repair the condition causing the I/O error. After the problem is repaired, run `fsck` and mount the file system again. Select `nodisable` if you want to implement the policy that most closely resembles the error handling policy of the previous VxFS release.

wdisable policy and mwdisable policy

If `wdisable` (write disable) or `mwdisable` (metadata-write disable) is selected, the file system is disabled or degraded, depending on the type of error encountered. Select `wdisable` or `mwdisable` for environments where read errors are more likely

to persist than write errors, such as when using non-redundant storage. `mdisable` is the default `ioerror` mount option for local mounts.

See the `mount_vxfs(1M)` manual page.

mdisable policy

If `mdisable` (metadata disable) is selected, the file system is disabled if a metadata read or write fails. However, the file system continues to operate if the failure is confined to data extents. `mdisable` is the default `ioerror` mount option for cluster mounts.

largefiles|nolargefiles mount option

VxFS supports files larger than two gigabytes. Files larger than 32 terabytes can be created only on 64-bit kernel operating systems and on a Veritas Volume Manager volume.

Note: Applications and utilities such as backup may experience problems if they are not aware of large files. In such a case, create your file system without large file capability.

See [“Creating a file system with large files”](#) on page 222.

See [“Mounting a file system with large files”](#) on page 223.

See [“Managing a file system with large files”](#) on page 223.

Creating a file system with large files

To create a file system with a file capability:

```
# mkfs -F vxfs -o largefiles special_device size
```

Specifying `largefiles` sets the `largefiles` flag. This lets the file system to hold files that are two gigabytes or larger. This is the default option.

To clear the flag and prevent large files from being created:

```
# mkfs -F vxfs -o nolargefiles special_device size
```

The `largefiles` flag is persistent and stored on disk.

Mounting a file system with large files

If a mount succeeds and `nolargefiles` is specified, the file system cannot contain or create any large files. If a mount succeeds and `largefiles` is specified, the file system may contain and create large files.

The `mount` command fails if the specified `largefiles|nolargefiles` option does not match the on-disk flag.

Because the `mount` command defaults to match the current setting of the on-disk flag if specified without the `largefiles` or `nolargefiles` option, the best practice is not to specify either option. After a file system is mounted, you can use the `fsadm` utility to change the large files option.

Managing a file system with large files

Managing a file system with large files includes the following tasks:

- Determining the current status of the large files flag
- Switching capabilities on a mounted file system
- Switching capabilities on an unmounted file system

To determine the current status of the `largefiles` flag, type either of the following commands:

```
# mkfs -F vxfs -m special_device
# fsadm -F vxfs mount_point | special_device
```

To switch capabilities on a mounted file system:

```
# fsadm -F vxfs -o [no]largefiles mount_point
```

To switch capabilities on an unmounted file system:

```
# fsadm -F vxfs -o [no]largefiles special_device
```

You cannot change a file system to `nolargefiles` if it contains large files.

See the `mount_vxfs(1M)`, `fsadm_vxfs(1M)`, and `mkfs_vxfs(1M)` manual pages.

cio mount option

The `cio` (Concurrent I/O) option specifies the file system to be mounted for concurrent reads and writes. Concurrent I/O is a licensed feature of VxFS. If `cio` is specified, but the license is not present, the `mount` command prints an error message and terminates the operation without mounting the file system. The `cio`

option cannot be disabled through a remount. To disable the `cio` option, the file system must be unmounted and mounted again without the `cio` option.

Note: There is no separate license required for the `cio` option, it is a part of the “standard” license.

mntlock mount option

The `mntlock` option prevents a file system from being unmounted by an application. This option is useful for applications that do not want the file systems that the applications are monitoring to be improperly unmounted by other applications or administrators.

The `mntunlock` option of the `umount` command reverses the `mntlock` option if you previously locked the file system.

ckptautomnt mount option

The `ckptautomnt` option enables the Storage Checkpoint visibility feature, which makes Storage Checkpoints easier to access.

See [“Storage Checkpoint visibility”](#) on page 575.

Combining mount command options

Although mount options can be combined arbitrarily, some combinations do not make sense. The following examples provide some common and reasonable mount option combinations.

To mount a desktop file system using options:

```
# mount -F vxfs -o log,mincache=closesync /dev/dsk/c1t3d0s1 /mnt  
  
# mount -F vxfs -o log,mincache=closesync /dev/dsk/c1t3d0s1 /mnt
```

This guarantees that when a file is closed, its data is synchronized to disk and cannot be lost. Thus, after an application has exited and its files are closed, no data is lost even if the system is immediately turned off.

To mount a temporary file system or to restore from backup:

```
# mount -F vxfs -o tmplog,convosync=delay,mincache=tmpcache \  
/dev/dsk/c1t3d0s1 /mnt
```

This combination might be used for a temporary file system where performance is more important than absolute data integrity. Any `O_SYNC` writes are performed as

delayed writes and delayed extending writes are not handled. This could result in a file that contains corrupted data if the system crashes. Any file written 30 seconds or so before a crash may contain corrupted data or be missing if this mount combination is in effect. However, such a file system does significantly less disk writes than a log file system, and should have significantly better performance, depending on the application.

To mount a file system for synchronous writes:

```
# mount -F vxfs -o log,convosync=dsync /dev/dsk/clt3d0s1 /mnt
```

This combination can be used to improve the performance of applications that perform `O_SYNC` writes, but only require data synchronous write semantics. Performance can be significantly improved if the file system is mounted using `convosync=dsync` without any loss of data integrity.

Example of mounting a file system

The following example mounts the file system `/dev/vx/dsk/fsvol/vol1` on the `/mnt1` directory with read/write access and delayed logging.

To mount the file system

- ◆ Mount the file system:

```
# mount -F vxfs -o delaylog /dev/vx/dsk/fsvol/vol1 /mnt1
```

Unmounting a file system

Use the `umount` command to unmount a currently mounted file system.

See the `umount_vxfs(1M)` manual page.

To unmount a file system

- ◆ Use the `umount` command to unmount a file system:

```
umount [-F vxfs] [generic_options] [-o [force]] {special|mount_point}
```

Specify the file system to be unmounted as a *mount_point* or *special*. *special* is the VxFS block special device on which the file system resides.

Example of unmounting a file system

The following are examples of unmounting file systems.

To unmount the file system `/dev/vx/dsk/fsvol/vol1`

- ◆ Unmount the file system:

```
# umount /dev/vx/dsk/fsvol/vol1
```

Resizing a file system

You can extend or shrink mounted VxFS file systems using the `fsadm` command. The size to which file system can be increased depends on the file system disk layout version. A file system using the Version 7 or later disk layout can be up to 8 exabytes in size. The size to which a Version 7 or later disk layout file system can be increased depends on the file system block size.

See [“About disk layouts”](#) on page 939.

See the `format(1M)` and `fsadm_vxfs(1M)` manual pages.

Extending a file system using fsadm

You can resize a file system by using the `fsadm` command.

To resize a VxFS file system

- ◆ Use the `fsadm` command to extend a VxFS file system:

```
fsadm [-F vxfs] [-b newsize] [-r rawdev] \  
mount_point
```

<code>vxfs</code>	The file system type.
<code>newsize</code>	The size to which the file system will increase. The default units is sectors, but you can specify <code>k</code> or <code>K</code> for kilobytes, <code>m</code> or <code>M</code> for megabytes, or <code>g</code> or <code>G</code> for gigabytes.
<code>mount_point</code>	The file system's mount point.
<code>-r rawdev</code>	Specifies the path name of the raw device if there is no entry in <code>/etc/vfstab</code> and <code>fsadm</code> cannot determine the raw device.

Examples of extending a file system

The following example extends a file system mounted at `/mnt1` to 22528 sectors.

To extend a file system to 22528 sectors

- ◆ Extend the VxFS file system mounted on `/mnt1` to 22528 sectors:

```
# fsadm -F vxfs -b 22528 /mnt1
```

The following example extends a file system mounted at `/mnt1` to 500 gigabytes.

To extend a file system to 500 gigabytes

- ◆ Extend the VxFS file system mounted on `/mnt1` to 500 gigabytes:

```
# fsadm -F vxfs -b 500g /mnt1
```

Shrinking a file system

You can decrease the size of the file system using `fsadm`, even while the file system is mounted.

Warning: After this operation, there is unused space at the end of the device. You can then resize the device, but be careful not to make the device smaller than the new size of the file system.

To decrease the size of a VxFS file system

- ◆ Use the `fsadm` command to decrease the size of a VxFS file system:

```
fsadm [-F vxfs] [-b newsize] [-r rawdev] mount_point
```

<code>vxfs</code>	The file system type.
<code>newsize</code>	The size to which the file system will shrink. The default units is sectors, but you can specify <code>k</code> or <code>K</code> for kilobytes, <code>m</code> or <code>M</code> for megabytes, or <code>g</code> or <code>G</code> for gigabytes.
<code>mount_point</code>	The file system's mount point.
<code>-r rawdev</code>	Specifies the path name of the raw device if there is no entry in <code>/etc/vfstab</code> and <code>fsadm</code> cannot determine the raw device.

Examples of shrinking a file system

The following example shrinks a VxFS file system mounted at `/mnt1` to 20480 sectors.

To shrink a file system to 20480 sectors

- ◆ Shrink a VxFS file system mounted at `/mnt1` to 20480 sectors:

```
# fsadm -F vxfs -b 20480 /mnt1
```

The following example shrinks a file system mounted at `/mnt1` to 450 gigabytes.

To shrink a file system to 450 gigabytes

- ◆ Shrink the VxFS file system mounted on `/mnt1` to 450 gigabytes:

```
# fsadm -F vxfs -b 450g /mnt1
```

Reorganizing a file system

You can reorganize or compact a fragmented file system using `fsadm`, even while the file system is mounted. This may help shrink a file system that could not previously be decreased.

To reorganize a VxFS file system

- ◆ Use the `fsadm` command to reorganize a VxFS file system:

```
fsadm [-F vxfs] [-e] [-d] [-E] [-D] [-H] [-r rawdev] mount_point
```

<code>vxfs</code>	The file system type.
<code>-d</code>	Reorders directory entries to put subdirectory entries first, then all other entries in decreasing order of time of last access. Also compacts directories to remove free space.
<code>-D</code>	Reports on directory fragmentation.
<code>-e</code>	Minimizes file system fragmentation. Files are reorganized to have the minimum number of extents.
<code>-E</code>	Reports on extent fragmentation.
<code>-H</code>	Displays the storage size in human friendly units (KB/MB/GB/TB/PB/EB), when used with the <code>-E</code> and <code>-D</code> options.
<code>mount_point</code>	The file system's mount point.
<code>-r rawdev</code>	Specifies the path name of the raw device if there is no entry in <code>/etc/vfstab</code> and <code>fsadm</code> cannot determine the raw device.

To perform free space defragmentation

- ◆ Use the `fsadm` command to perform free space defragmentation of a VxFS file system:

```
fsadm [-F vxfs] [-C] mount_point
```

<code>vxfs</code>	The file system type.
<code>-C</code>	Minimizes file system free space fragmentation. This attempts to generate bigger chunks of free space in the device.
<code>mount_point</code>	The file system's mount point.

Example of reorganizing a file system

The following example reorganizes the file system mounted at `/mnt1`.

To reorganize a VxFS file system

- ◆ Reorganize the VxFS file system mounted at `/mnt1`:

```
# fsadm -F vxfs -EeDd /mnt1
```

Example of running free space defragmentation

The following example minimizes the free space fragmentation of the file system mounted at `/mnt1`.

To run free space defragmentation

- ◆ Minimize the free space of the the VxFS file system mounted at `/mnt1`:

```
# fsadm -F vxfs -C /mnt1
```

Displaying information on mounted file systems

Use the `mount` command to display a list of currently mounted file systems.

See the `mount(1M)` and `mount_vxfs(1M)` manual pages.

To view the status of mounted file systems

- ◆ Use the `mount` command to view the status of mounted file systems:

```
mount -v
```

This shows the file system type and `mount` options for all mounted file systems. The `-v` option specifies verbose mode.

Example of displaying information on mounted file systems

The following example shows the result of invoking the `mount` command without options.

To display information on mounted file systems

- ◆ Invoke the `mount` command without options:

```
# mount
/ on /dev/root read/write/setuid on Mon Aug 22 16:58:24 2011
/proc on /proc read/write on Mon Aug 22 16:58:25 2011
/dev/fd on /dev/fd read/write on Mon Aug 22 16:58:26 2011
/tmp on /tmp read/write on Mon Aug 22 16:59:33 2011
/var/tmp on /var/tmp read/write on Mon Aug 22 16:59:34 2011
```

Identifying file system types

Use the `fstyp` command to determine the file system type for a specified file system. This is useful when a file system was created elsewhere and you want to know its type.

See the `fstyp(1M)` and `fstyp_vxfs(1M)` manual pages.

To determine a file system's type

- ◆ Use the `fstyp` command to determine a file system's type:

```
fstyp -v special
```

special The character (raw) device.

`-v` Specifies verbose mode.

Example of determining a file system's type

The following example uses the `fstyp` command to determine a the file system type of the `/dev/vx/dsk/fsvol/vol1` device.

To determine the file system's type

- ◆ Use the `fstyp` command to determine the file system type of the device `/dev/vx/dsk/fsvol/vol1`:

```
# fstyp -v /dev/vx/dsk/fsvol/vol1
```

The output indicates that the file system type is vxfs, and displays file system information similar to the following:

```
vxfs
magic a501fcf5 version 7 ctime Tue Jun 25 18:29:39 2003
logstart 17 logend 1040
bsize 1024 size 1048576 dsize 1047255 ninode 0 nau 8
defiextsize 64 ilbsize 0 immedlen 96 ndaddr 10
aufirst 1049 emap 2 imap 0 iextop 0 istart 0
bstart 34 femap 1051 fimap 0 fiextop 0 fistart 0 fbstart
```

Monitoring free space

In general, VxFS works best if the percentage of free space in the file system does not get below 10 percent. This is because file systems with 10 percent or more free space have less fragmentation and better extent allocation. Regular use of the `df` command to monitor free space is desirable.

See the `df_vxfs(1M)` manual page.

Full file systems may have an adverse effect on file system performance. Full file systems should therefore have some files removed, or should be expanded.

See the `fsadm_vxfs(1M)` manual page.

Monitoring fragmentation

Fragmentation reduces performance and availability. Regular use of `fsadm's` fragmentation reporting and reorganization facilities is therefore advisable.

The easiest way to ensure that fragmentation does not become a problem is to schedule regular defragmentation runs using the `cron` command.

Defragmentation scheduling should range from weekly (for frequently used file systems) to monthly (for infrequently used file systems). Extent fragmentation should be monitored with `fsadm` command.

To determine the degree of fragmentation, use the following factors:

- Percentage of free space in extents of less than 8 blocks in length
- Percentage of free space in extents of less than 64 blocks in length
- Percentage of free space in extents of length 64 blocks or greater

An unfragmented file system has the following characteristics:

- Less than 1 percent of free space in extents of less than 8 blocks in length
- Less than 5 percent of free space in extents of less than 64 blocks in length
- More than 5 percent of the total file system size available as free extents in lengths of 64 or more blocks

A badly fragmented file system has one or more of the following characteristics:

- Greater than 5 percent of free space in extents of less than 8 blocks in length
- More than 50 percent of free space in extents of less than 64 blocks in length
- Less than 5 percent of the total file system size available as free extents in lengths of 64 or more blocks

Fragmentation can also be determined based on the fragmentation index. Two types of indices are generated by the `fsadm` command: the file fragmentation index and the free space fragmentation index. Both of these indices range between 0 and 100, and give an idea about the level of file fragmentation and free space fragmentation, respectively. A value of 0 for the fragmentation index means that the file system has no fragmentation, and a value of 100 means that the file system has the highest level of fragmentation. Based on the index, you should use the appropriate defragmentation option with the `fsadm` command. For example if the file fragmentation index is high, the `fsadm` command should be run with the `-e` option. If the free space fragmentation index is high, the `fsadm` command should be run with `-C` option. When the `fsadm` command is run with the `-e` option, internally it performs free space defragmentation before performing file defragmentation.

The optimal period for scheduling of extent reorganization runs can be determined by choosing a reasonable interval, scheduling `fsadm` runs at the initial interval, and running the extent fragmentation report feature of `fsadm` before and after the reorganization.

The “before” result is the degree of fragmentation prior to the reorganization. If the degree of fragmentation is approaching the figures for bad fragmentation, reduce

the interval between `fsadm` runs. If the degree of fragmentation is low, increase the interval between `fsadm` runs.

The “after” result is an indication of how well the reorganizer has performed. The degree of fragmentation should be close to the characteristics of an unfragmented file system. If not, it may be a good idea to resize the file system; full file systems tend to fragment and are difficult to defragment. It is also possible that the reorganization is not being performed at a time during which the file system in question is relatively idle.

Directory reorganization is not nearly as critical as extent reorganization, but regular directory reorganization improves performance. It is advisable to schedule directory reorganization for file systems when the extent reorganization is scheduled. The following is a sample script that is run periodically at 3:00 A.M. from `cron` for a number of file systems:

```
outfile=/usr/spool/fsadm/out.`/bin/date +%m%d`
for i in /home /home2 /project /db
do
    /bin/echo "Reorganizing $i"
    /bin/timex fsadm -F vxfs -e -E -s $i
    /bin/timex fsadm -F vxfs -s -d -D $i
done > $outfile 2>&1
```

Thin Reclamation

Veritas File System (VxFS) supports reclamation of free storage on a Thin Storage LUN.

You reclaim free storage using the `fsadm` command or the `vxfs_ts_reclaim` API. You can perform the default reclamation or aggressive reclamation. If you used a file system for a long time and must perform reclamation on the file system, Symantec recommends that you run aggressive reclamation. Aggressive reclamation compacts the allocated blocks, which creates larger free blocks that can potentially be reclaimed.

You can specify the following thin reclamation options with the `fsadm` command:

<code>-o aggressive</code>	Initiates Thin Storage aggressive reclamation.
<code>-o analyze analyze</code>	Initiates the analyze reclaim option.
<code>-o auto</code>	Initiates the auto reclaim option.

- P Performs multi-threaded Thin Storage Reclamation. By default, the `fsadm` command performs single-threaded Thin Storage Reclamation. To use multi-threaded Thin Storage Reclamation, the array must support multiple concurrent reclaim operations.
- R Performs reclamation of free storage to the Thin Storage LUN on a VxFS file system .

See the `fsadm_vxfs(1M)` and `vxfs_ts_reclaim(3)` manual pages.

Thin Reclamation is only supported on file systems mounted on a VxVM volume.

The following example performs default reclamation of free storage to the Thin Storage LUN on a VxFS file system mounted at `/mnt1`:

```
# fsadm -R /mnt1
```

The following example performs aggressive reclamation of free storage to the Thin Storage LUN on a VxFS file system mounted at `/mnt1`:

```
# fsadm -R -o aggressive /mnt1
```

After performing the reclaim operation, you can verify that the storage was reclaimed using the `vxdisk -o thin list` command.

Veritas File System also supports reclamation of a portion of the file system using the `vxfs_ts_reclaim()` API.

See the *Veritas File System Programmer's Reference Guide*.

Note: Thin Reclamation is a slow process and may take several hours to complete, depending on the file system size. Thin Reclamation is not guaranteed to reclaim 100% of the free space.

You can track the progress of the Thin Reclamation process by using the `vxtask list` command when using the Veritas Volume Manager (VxVM) command `vxdisk reclaim`.

See the `vxtask(1M)` and `vxdisk(1M)` manual pages.

You can administer Thin Reclamation using VxVM commands.

See [“Thin Reclamation of a disk, a disk group, or an enclosure”](#) on page 684.

Extent attributes

This chapter includes the following topics:

- [About extent attributes](#)
- [Commands related to extent attributes](#)

About extent attributes

Veritas Storage Foundation Cluster File System High Availability (VxFS) allocates disk space to files in groups of one or more adjacent blocks called extents. VxFS defines an application interface that allows programs to control various aspects of the extent allocation for a given file. The extent allocation policies associated with a file are referred to as extent attributes.

The VxFS `getext` and `setext` commands let you view or manipulate file extent attributes. In addition, the `vxdump`, `vxrestore`, `mv_vxfs`, `cp_vxfs`, and `cpio_vxfs` commands preserve extent attributes when a file is backed up, moved, copied, or archived.

The two basic extent attributes associated with a file are its reservation and its fixed extent size. You can preallocate space to the file by manipulating a file's reservation, or override the default allocation policy of the file system by setting a fixed extent size.

See [“Reservation: preallocating space to a file”](#) on page 236.

See [“Fixed extent size”](#) on page 236.

Other policies determine the way these attributes are expressed during the allocation process.

You can specify the following criteria:

- The space reserved for a file must be contiguous
- No allocations will be made for a file beyond the current reservation

- An unused reservation will be released when the file is closed
- Space will be allocated, but no reservation will be assigned
- The file size will be changed to incorporate the allocated space immediately

Some of the extent attributes are persistent and become part of the on-disk information about the file, while other attributes are temporary and are lost after the file is closed or the system is rebooted. The persistent attributes are similar to the file's permissions and are written in the inode for the file. When a file is copied, moved, or archived, only the persistent attributes of the source file are preserved in the new file.

See [“Other extent attribute controls”](#) on page 237.

In general, the user will only set extent attributes for reservation. Many of the attributes are designed for applications that are tuned to a particular pattern of I/O or disk alignment.

See the `setext(1)` manual page.

See [“About Veritas File System I/O”](#) on page 485.

Reservation: preallocating space to a file

VxFS makes it possible to preallocate space to a file at the time of the request rather than when data is written into the file. This space cannot be allocated to other files in the file system. VxFS prevents any unexpected out-of-space condition on the file system by ensuring that a file's required space will be associated with the file before it is required.

A persistent reservation is not released when a file is truncated. The reservation must be cleared or the file must be removed to free the reserved space.

Fixed extent size

The VxFS default allocation policy uses a variety of methods to determine how to make an allocation to a file when a write requires additional space. The policy attempts to balance the two goals of optimum I/O performance through large allocations and minimal file system fragmentation. VxFS accomplishes these goals by allocating from space available in the file system that best fits the data.

Setting a fixed extent size overrides the default allocation policies for a file and always serves as a persistent attribute. Be careful to choose an extent size appropriate to the application when using fixed extents. An advantage of the VxFS extent-based allocation policies is that they rarely use indirect blocks compared to block based file systems; VxFS eliminates many instances of disk access that stem from indirect references. However, a small extent size can eliminate this advantage.

Files with large extents tend to be more contiguous and have better I/O characteristics. However, the overall performance of the file system degrades because the unused space fragments free space by breaking large extents into smaller pieces. By erring on the side of minimizing fragmentation for the file system, files may become so non-contiguous that their I/O characteristics would degrade.

Fixed extent sizes are particularly appropriate in the following situations:

- If a file is large and sparse and its write size is fixed, a fixed extent size that is a multiple of the write size can minimize space wasted by blocks that do not contain user data as a result of misalignment of write and extent sizes. The default extent size for a sparse file is 8K.
- If a file is large and contiguous, a large fixed extent size can minimize the number of extents in the file.

Custom applications may also use fixed extent sizes for specific reasons, such as the need to align extents to cylinder or striping boundaries on disk.

How the fixed extent size works with the shared extents

Veritas File System (VxFS) allows the user to set the fixed extent size option on a file that controls the minimum allocation size of the file. If a file has shared extents that must be unshared, the allocation that is done as a part of the unshare operation ignores the fixed extent size option that is set on the file. The allocation size during the unshare operation, is dependent on the size of the write operation on the shared region.

Other extent attribute controls

The auxiliary controls on extent attributes determine the following conditions:

- Whether allocations are aligned
See [“Extent attribute alignment”](#) on page 238.
- Whether allocations are contiguous
See [“Extent attribute contiguity”](#) on page 238.
- Whether the file can be written beyond its reservation
See [“Write operations beyond extent attribute reservation”](#) on page 238.
- Whether an unused reservation is released when the file is closed
See [“Extent attribute reservation trimming”](#) on page 238.
- Whether the reservation is a persistent attribute of the file
See [“Extent attribute reservation persistence”](#) on page 238.
- When the space reserved for a file will actually become part of the file
See [“Including an extent attribute reservation in the file”](#) on page 238.

Extent attribute alignment

Specific alignment restrictions coordinate a file's allocations with a particular I/O pattern or disk alignment. Alignment can only be specified if a fixed extent size has also been set. Setting alignment restrictions on allocations is best left to well-designed applications.

See the `setext(1)` manual page.

See [“About Veritas File System I/O”](#) on page 485.

Extent attribute contiguity

A reservation request can specify that its allocation remain contiguous (all one extent). Maximum contiguity of a file optimizes its I/O characteristics.

Note: Fixed extent sizes or alignment cause a file system to return an error message reporting insufficient space if no suitably sized (or aligned) extent is available. This can happen even if the file system has sufficient free space and the fixed extent size is large.

Write operations beyond extent attribute reservation

A reservation request can specify that no allocations can take place after a write operation fills the last available block in the reservation. This request can be used a way similar to the function of the `ulimit` command to prevent a file's uncontrolled growth.

Extent attribute reservation trimming

A reservation request can specify that any unused reservation be released when the file is closed. The file is not completely closed until all processes open against the file have closed it.

Extent attribute reservation persistence

A reservation request can ensure that the reservation does not become a persistent attribute of the file. The unused reservation is discarded when the file is closed.

Including an extent attribute reservation in the file

A reservation request can make sure the size of the file is adjusted to include the reservation. Normally, the space of the reservation is not included in the file until an extending write operation requires it. A reservation that immediately changes the file size can generate large temporary files. Unlike a `ftruncate` operation that

increases the size of a file, this type of reservation does not perform zeroing of the blocks included in the file and limits this facility to users with appropriate privileges. The data that appears in the file may have been previously contained in another file. For users who do not have the appropriate privileges, there is a variant request that prevents such users from viewing uninitialized data.

Commands related to extent attributes

The VxFS commands for manipulating extent attributes are `setext` and `getext`; they allow the user to set up files with a given set of extent attributes or view any attributes that are already associated with a file.

See the `setext(1)` and `getext(1)` manual pages.

The VxFS-specific commands `vxdump`, `vxrestore`, `mv_vxfs`, `cp_vxfs`, and `cpio_vxfs` preserve extent attributes when backing up, restoring, moving, or copying files. Make sure to modify your *PATH* when using the VxFS versions of `mv`, `cp`, and `cpio`.

Most of these commands include a command line option (`-e`) for maintaining extent attributes on files. This option specifies dealing with a VxFS file that has extent attribute information including reserved space, a fixed extent size, and extent alignment. The extent attribute information may be lost if the destination file system does not support extent attributes, has a different block size than the source file system, or lacks free extents appropriate to satisfy the extent attribute requirements.

The `-e` option takes any of the following keywords as an argument:

<code>warn</code>	Issues a warning message if extent attribute information cannot be maintained (the default)
<code>force</code>	Fails the copy if extent attribute information cannot be maintained
<code>ignore</code>	Ignores extent attribute information entirely

Example of setting an extent attribute

The following example creates a file named `file1` and preallocates 2 GB of disk space for the file.

To set an extent attribute

- 1 Create the file `file1`:

```
# touch file1
```

- 2 Preallocate 2 GB of disk space for the file `file1`:

```
# setext -F vxfs -r 2g -f chgsize file1
```

Since the example specifies the `-f chgsize` option, VxFS immediately incorporates the reservation into the file and updates the file's inode with size and block count information that is increased to include the reserved space. Only users with root privileges can use the `-f chgsize` option.

Example of getting an extent attribute

The following example gets the extent attribute information of a file named `file1`.

To get an extent attribute's information

- ◆ Get the extent attribute information for the file `file1`:

```
# getext -F vxfs file1
file1: Bsize 1024 Reserve 36 Extent Size 3 align noextend
```

The file `file1` has a block size of 1024 bytes, 36 blocks reserved, a fixed extent size of 3 blocks, and all extents aligned to 3 block boundaries. The file size cannot be increased after the current reservation is exhausted. Reservations and fixed extent sizes are allocated in units of the file system block size.

Failure to preserve extent attributes

Whenever a file is copied, moved, or archived using commands that preserve extent attributes, there is the possibility of losing the attributes.

Such a failure might occur for one of the following reasons:

- The file system receiving a copied, moved, or restored file from an archive is not a VxFS type. Since other file system types do not support the extent attributes of the VxFS file system, the attributes of the source file are lost during the migration.
- The file system receiving a copied, moved, or restored file is a VxFS type but does not have enough free space to satisfy the extent attributes. For example, consider a 50K file and a reservation of 1 MB. If the target file system has 500K free, it could easily hold the file but fail to satisfy the reservation.

- The file system receiving a copied, moved, or restored file from an archive is a VxFS type but the different block sizes of the source and target file system make extent attributes impossible to maintain. For example, consider a source file system of block size 1024, a target file system of block size 4096, and a file that has a fixed extent size of 3 blocks (3072 bytes). This fixed extent size adapts to the source file system but cannot translate onto the target file system.

The same source and target file systems in the preceding example with a file carrying a fixed extent size of 4 could preserve the attribute; a 4 block (4096 byte) extent on the source file system would translate into a 1 block extent on the target.

On a system with mixed block sizes, a copy, move, or restoration operation may or may not succeed in preserving attributes. It is recommended that the same block size be used for all file systems on a given system.

Administering multi-pathing with DMP

- [Chapter 11. Administering Dynamic Multi-Pathing](#)
- [Chapter 12. Dynamic reconfiguration of devices](#)
- [Chapter 13. Managing devices](#)
- [Chapter 14. Event monitoring](#)

Administering Dynamic Multi-Pathing

This chapter includes the following topics:

- [Discovering and configuring newly added disk devices](#)
- [Making devices invisible to VxVM](#)
- [Making devices visible to VxVM](#)
- [About enabling and disabling I/O for controllers and storage processors](#)
- [About displaying Veritas Dynamic Multi-Pathing \(DMP\) information](#)
- [Displaying the paths to a disk](#)
- [Administering DMP using vxdkmpadm](#)
- [DMP coexistence with native multi-pathing](#)

Discovering and configuring newly added disk devices

The `vxdiskconfig` utility scans and configures new disk devices attached to the host, disk devices that become online, or fibre channel devices that are zoned to host bus adapters connected to this host. The command calls platform specific interfaces to configure new disk devices and brings them under control of the operating system. It scans for disks that were added since SFCFSA's configuration daemon was last started. These disks are then dynamically configured and recognized by SFCFSA.

`vxdiskconfig` should be used whenever disks are physically connected to the host or when fibre channel devices are zoned to the host.

`vxdiskconfig` calls `vxctl enable` to rebuild volume device node directories and update the DMP internal database to reflect the new state of the system.

You can also use the `vxdisk scandisks` command to scan devices in the operating system device tree, and to initiate dynamic reconfiguration of multipathed disks.

If you want SFCFSA to scan only for new devices that have been added to the system, and not for devices that have been enabled or disabled, specify the `-f` option to either of the commands, as shown here:

```
# vxctl -f enable
# vxdisk -f scandisks
```

However, a complete scan is initiated if the system configuration has been modified by changes to:

- Installed array support libraries.
- The list of devices that are excluded from use by VxVM.
- DISKS (JBOD), SCSI3, or foreign device definitions.

See the `vxctl(1M)` manual page.

See the `vxdisk(1M)` manual page.

Partial device discovery

Dynamic Multi-Pathing (DMP) supports partial device discovery where you can include or exclude paths to a physical disk from the discovery process.

The `vxdisk scandisks` command rescans the devices in the OS device tree and triggers a DMP reconfiguration. You can specify parameters to `vxdisk scandisks` to implement partial device discovery. For example, this command makes SFCFSA discover newly added devices that were unknown to it earlier:

```
# vxdisk scandisks new
```

The next example discovers fabric devices:

```
# vxdisk scandisks fabric
```

The above command discovers devices with the characteristic `DDI_NT_FABRIC` property set on them.

The following command scans for the devices `c1t1d0` and `c2t2d0`:

```
# vxdisk scandisks device=c1t1d0,c2t2d0
```


Alternatively, you can specify a `!` prefix character to indicate that you want to scan for all devices except those that are listed.

Note: The `!` character is a special character in some shells. The following examples show how to escape it in a bash shell.

```
# vxdisk scandisks \!device=c1t1d0,c2t2d0
```

You can also scan for devices that are connected (or not connected) to a list of logical or physical controllers. For example, this command discovers and configures all devices except those that are connected to the specified logical controllers:

```
# vxdisk scandisks \!ctlr=c1,c2
```

The next command discovers devices that are connected to the specified physical controller:

```
# vxdisk scandisks pctlr=/pci@1f,4000/scsi@3/
```

The items in a list of physical controllers are separated by `+` characters.

You can use the command `vxmpadm getctlr all` to obtain a list of physical controllers.

You should specify only one selection argument to the `vxdisk scandisks` command. Specifying multiple options results in an error.

See the `vxdisk(1M)` manual page.

Discovering disks and dynamically adding disk arrays

DMP uses array support libraries (ASLs) to provide array-specific support for multi-pathing. An array support library (ASL) is a dynamically loadable shared library (plug-in for DDL). The ASL implements hardware-specific logic to discover device attributes during device discovery. DMP provides the device discovery layer (DDL) to determine which ASLs should be associated to each disk array.

In some cases, DMP can also provide basic multi-pathing and failover functionality by treating LUNs as disks (JBODs).

How DMP claims devices

For fully optimized support of any array and for support of more complicated array types, DMP requires the use of array-specific array support libraries (ASLs), possibly coupled with array policy modules (APMs). ASLs and APMs effectively are array-specific plugins that allow close tie-in of DMP with any specific array model.

See the Hardware Compatibility List for the complete list of supported arrays.

<http://www.symantec.com/docs/TECH170013>

During device discovery, the DDL checks the installed ASL for each device to find which ASL claims the device.

If no ASL is found to claim the device, the DDL checks for a corresponding JBOD definition. You can add JBOD definitions for unsupported arrays to enable DMP to provide multi-pathing for the array. If a JBOD definition is found, the DDL claims the devices in the DISKS category, which adds the LUNs to the list of JBOD (physical disk) devices used by DMP. If the JBOD definition includes a cabinet number, DDL uses the cabinet number to group the LUNs into enclosures.

See “[Adding unsupported disk arrays to the DISKS category](#)” on page 258.

DMP can provide basic multi-pathing to ALUA-compliant arrays even if there is no ASL or JBOD definition. DDL claims the LUNs as part of the aluadisk enclosure. The array type is shown as ALUA. Adding a JBOD definition also enables you to group the LUNs into enclosures.

Disk categories

Disk arrays that have been certified for use with Veritas Storage Foundation Cluster File System High Availability are supported by an array support library (ASL), and are categorized by the vendor ID string that is returned by the disks (for example, “HITACHI”).

Disks in JBODs which are capable of being multipathed by DMP, are placed in the `DISKS` category. Disks in unsupported arrays can also be placed in the `DISKS` category.

See “[Adding unsupported disk arrays to the DISKS category](#)” on page 258.

Disks in JBODs that do not fall into any supported category, and which are not capable of being multipathed by DMP are placed in the `OTHER_DISKS` category.

Adding support for a new disk array

You can dynamically add support for a new type of disk array. The support comes in the form of Array Support Libraries (ASLs) that are developed by Symantec. Symantec provides support for new disk arrays through updates to the `VRTSaslapm` package. To determine if an updated `VRTSaslapm` package is available for download, refer to the hardware compatibility list tech note. The hardware compatibility list provides a link to the latest package for download and instructions for installing the `VRTSaslapm` package. You can upgrade the `VRTSaslapm` package while the system is online; you do not need to stop the applications.

To access the hardware compatibility list, go to the following URL:

<http://www.symantec.com/docs/TECH170013>

Each VRTSaslapm package is specific for the Storage Foundation Cluster File System High Availability version. Be sure to install the VRTSaslapm package that supports the installed version of Storage Foundation Cluster File System High Availability.

The new disk array does not need to be already connected to the system when the VRTSaslapm package is installed. If any of the disks in the new disk array are subsequently connected, and if vxconfigd is running, vxconfigd immediately invokes the Device Discovery function and includes the new disks in the VxVM device list.

See “Adding new LUNs dynamically to a new target ID” on page 314.

If you need to remove the latest VRTSaslapm package, you can revert to the previously installed version.

For the detailed procedure, refer to the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide*.

Enabling discovery of new disk arrays

The vxctl enable command scans all of the disk devices and their attributes, updates the SFCFSA device list, and reconfigures DMP with the new device database. There is no need to reboot the host.

Warning: This command ensures that Dynamic Multi-Pathing is set up correctly for the array. Otherwise, VxVM treats the independent paths to the disks as separate devices, which can result in data corruption.

To enable discovery of a new disk array

- ◆ Type the following command:

```
# vxctl enable
```

Third-party driver coexistence

The third-party driver (TPD) coexistence feature of SFCFSA allows I/O that is controlled by some third-party multi-pathing drivers to bypass DMP while retaining the monitoring capabilities of DMP. If a suitable ASL is available and installed, devices that use TPDs can be discovered without requiring you to set up a specification file, or to run a special command. In previous releases, VxVM only

supported TPD coexistence if the code of the third-party driver was intrusively modified. Now, the TPD coexistence feature maintains backward compatibility with such methods, but it also permits coexistence without requiring any change in a third-party multi-pathing driver.

See [“Changing device naming for TPD-controlled enclosures”](#) on page 330.

See [“Displaying DMP path information for devices under third-party driver control”](#) on page 279.

Autodiscovery of EMC Symmetrix arrays

In VxVM 4.0, there were two possible ways to configure EMC Symmetrix arrays:

- With EMC PowerPath installed, EMC Symmetrix arrays could be configured as foreign devices.
See [“Foreign devices”](#) on page 261.
- Without EMC PowerPath installed, DMP could be used to perform multi-pathing.

On upgrading a system to VxVM 4.1 or later release, existing EMC PowerPath devices can be discovered by DDL, and configured into DMP as autoconfigured disks with DMP nodes, even if PowerPath is being used to perform multi-pathing. There is no need to configure such arrays as foreign devices.

[Table 11-1](#) shows the scenarios for using DMP with PowerPath.

The ASLs are all included in the ASL-APM package, which is installed when you install Storage Foundation products.

Table 11-1 Scenarios for using DMP with PowerPath

PowerPath	DMP	Array configuration mode
Installed.	The <code>libvxpp</code> ASL handles EMC Symmetrix arrays and DGC CLARiON claiming internally. PowerPath handles failover.	EMC Symmetrix - Any DGC CLARiON - Active/Passive (A/P), Active/Passive in Explicit Failover mode (A/P-F) and ALUA Explicit failover
Not installed; the array is EMC Symmetrix.	DMP handles multi-pathing. The ASL name is <code>libvxemc</code> .	Active/Active

Table 11-1 Scenarios for using DMP with PowerPath (*continued*)

PowerPath	DMP	Array configuration mode
Not installed; the array is DGC CLARiioN (CXn00).	DMP handles multi-pathing. The ASL name is libvxCLARiioN.	Active/Passive (A/P), Active/Passive in Explicit Failover mode (A/P-F) and ALUA

If any EMCpower disks are configured as foreign disks, use the `vxddladm rmforeign` command to remove the foreign definitions, as shown in this example:

```
# vxddladm rmforeign blockpath=/dev/dsk/emcpower10 \
  charpath=/dev/rdisk/emcpower10
```

To allow DMP to receive correct inquiry data, the Common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

How to administer the Device Discovery Layer

The Device Discovery Layer (DDL) allows dynamic addition of disk arrays. DDL discovers disks and their attributes that are required for SFCFSA operations.

The DDL is administered using the `vxddladm` utility to perform the following tasks:

- List the hierarchy of all the devices discovered by DDL including iSCSI devices.
See [“Listing all the devices including iSCSI”](#) on page 250.
- List all the Host Bus Adapters including iSCSI
See [“Listing all the Host Bus Adapters including iSCSI”](#) on page 251.
- List the ports configured on a Host Bus Adapter
See [“Listing the ports configured on a Host Bus Adapter”](#) on page 252.
- List the targets configured from a Host Bus Adapter
See [“Listing the targets configured from a Host Bus Adapter or a port ”](#) on page 252.
- List the devices configured from a Host Bus Adapter
See [“Listing the devices configured from a Host Bus Adapter and target ”](#) on page 253.
- Get or set the iSCSI operational parameters
See [“Getting or setting the iSCSI operational parameters”](#) on page 254.
- List the types of arrays that are supported.
See [“Listing all supported disk arrays”](#) on page 256.

- Excluding support for an array from DDL.
See [“Excluding support for a disk array library”](#) on page 256.
- Re-including support for an array to DDL.
See [“Re-including support for an excluded disk array library”](#) on page 257.
- List information about excluded disk arrays.
See [“Listing excluded disk arrays”](#) on page 257.
- List disks that are supported in the `DISKS` (JBOD) category.
See [“Listing supported disks in the DISKS category”](#) on page 257.
- List information about supported array libraries.
See [“Displaying details about a supported array library”](#) on page 256.
- Add disks from different vendors to the `DISKS` category.
See [“Adding unsupported disk arrays to the DISKS category”](#) on page 258.
- Remove disks from the `DISKS` category.
See [“Removing disks from the DISKS category”](#) on page 260.
- Add disks as foreign devices.
See [“Foreign devices”](#) on page 261.

The referenced sections explain these tasks in more detail.

See the `vxddladm(1M)` manual page.

Listing all the devices including iSCSI

You can display the hierarchy of all the devices discovered by DDL, including iSCSI devices.

To list all the devices including iSCSI

- ◆ Type the following command:

```
# vxddladm list
```

The following is a sample output:

```
HBA c2 (20:00:00:E0:8B:19:77:BE)
    Port c2_p0 (50:0A:09:80:85:84:9D:84)
        Target c2_p0_t0 (50:0A:09:81:85:84:9D:84)
            LUN c2t0d0s2
. . .
HBA c3 (iqn.1986-03.com.sun:01:0003ba8ed1b5.45220f80)
    Port c3_p0 (10.216.130.10:3260)
        Target c3_p0_t0 (iqn.1992-08.com.netapp:sn.84188548)
            LUN c3t0d0s2
            LUN c3t0d1s2
        Target c3_t1 (iqn.1992-08.com.netapp:sn.84190939)
. . .
```

Listing all the Host Bus Adapters including iSCSI

You can obtain information about all the Host Bus Adapters configured on the system, including iSCSI adapters. This includes the following information:

Driver	Driver controlling the HBA.
Firmware	Firmware version.
Discovery	The discovery method employed for the targets.
State	Whether the device is Online or Offline.
Address	The hardware address.

To list all the Host Bus Adapters including iSCSI

- ◆ Use the following command to list all of the HBAs, including iSCSI devices, configured on the system:

```
# vxddladm list hbas
```

Listing the ports configured on a Host Bus Adapter

You can obtain information about all the ports configured on an HBA. The display includes the following information:

HBA-ID	The parent HBA.
State	Whether the device is Online or Offline.
Address	The hardware address.

To list the ports configured on a Host Bus Adapter

- ◆ Use the following command to obtain the ports configured on an HBA:

```
# vxddladm list ports
```

PortID	HBA-ID	State	Address

c2_p0	c2	Online	50:0A:09:80:85:84:9D:84
c3_p0	c3	Online	10.216.130.10:3260

Listing the targets configured from a Host Bus Adapter or a port

You can obtain information about all the targets configured from a Host Bus Adapter or a port. This includes the following information:

Alias	The alias name, if available.
HBA-ID	Parent HBA or port.
State	Whether the device is Online or Offline.
Address	The hardware address.

To list the targets

- ◆ To list all of the targets, use the following command:

```
# vxddladm list targets
```

The following is a sample output:

TgtID	Alias	HBA-ID	State	Address
c2_p0_t0	-	c2	Online	50:0A:09:80:85:84:9D:84
c3_p0_t1	-	c3	Online	iqn.1992-08.com.netapp:sn.84190939

To list the targets configured from a Host Bus Adapter or port

- ◆ You can filter based on a HBA or port, using the following command:

```
# vxddladm list targets [hba=hba_name|port=port_name]
```

For example, to obtain the targets configured from the specified HBA:

```
# vxddladm list targets hba=c2
```

TgtID	Alias	HBA-ID	State	Address
c2_p0_t0	-	c2	Online	50:0A:09:80:85:84:9D:84

Listing the devices configured from a Host Bus Adapter and target

You can obtain information about all the devices configured from a Host Bus Adapter. This includes the following information:

Device	The device name.
Target-ID	The parent target.
State	Whether the device is Online or Offline.
DDL status	Whether the device is claimed by DDL. If claimed, the output also displays the ASL name.

To list the devices configured from a Host Bus Adapter

- ◆ To obtain the devices configured, use the following command:

```
# vxddladm list devices
```

Device	Target-ID	State	DDL status (ASL)
-----	-----	-----	-----
c2t0d2s2	c2_p0_t0	Online	CLAIMED (libvxemc.so)
c3t1d2s2	c3_p0_t1	Online	SKIPPED (libvxemc.so)
c4t1d2s2	c4_p0_t1	Offline	ERROR
c4t1d2s2	c4_p0_t2	Online	EXCLUDED
c4t5d2s2	c4_p0_t5	Offline	MASKED

To list the devices configured from a Host Bus Adapter and target

- ◆ To obtain the devices configured from a particular HBA and target, use the following command:

```
# vxddladm list devices target=target_name
```

Getting or setting the iSCSI operational parameters

DDL provides an interface to set and display certain parameters that affect the performance of the iSCSI device path. However, the underlying OS framework must support the ability to set these values. The `vxddladm set` command returns an error if the OS support is not available.

Table 11-2 Parameters for iSCSI devices

Parameter	Default value	Minimum value	Maximum value
DataPDUIInOrder	yes	no	yes
DataSequenceInOrder	yes	no	yes
DefaultTime2Retain	20	0	3600
DefaultTime2Wait	2	0	3600
ErrorRecoveryLevel	0	0	2
FirstBurstLength	65535	512	16777215
InitialR2T	yes	no	yes
ImmediateData	yes	no	yes

Table 11-2 Parameters for iSCSI devices (*continued*)

Parameter	Default value	Minimum value	Maximum value
MaxBurstLength	262144	512	16777215
MaxConnections	1	1	65535
MaxOutStandingR2T	1	1	65535
MaxRecvDataSegmentLength	8182	512	16777215

To get the iSCSI operational parameters on the initiator for a specific iSCSI target

- ◆ Type the following commands:

```
# vxddladm getiscsi target=tgt-id {all | parameter}
```

You can use this command to obtain all the iSCSI operational parameters.

```
# vxddladm getiscsi target=c2_p2_t0
```

The following is a sample output:

PARAMETER	CURRENT	DEFAULT	MIN	MAX
-----	-----	-----	-----	-----
DataPDUIInOrder	yes	yes	no	yes
DataSequenceInOrder	yes	yes	no	yes
DefaultTime2Retain	20	20	0	3600
DefaultTime2Wait	2	2	0	3600
ErrorRecoveryLevel	0	0	0	2
FirstBurstLength	65535	65535	512	16777215
InitialR2T	yes	yes	no	yes
ImmediateData	yes	yes	no	yes
MaxBurstLength	262144	262144	512	16777215
MaxConnections	1	1	1	65535
MaxOutStandingR2T	1	1	1	65535
MaxRecvDataSegmentLength	8192	8182	512	16777215

To set the iSCSI operational parameters on the initiator for a specific iSCSI target

- ◆ Type the following command:

```
# vxddladm setiscsi target=tgt-id parameter=value
```

Listing all supported disk arrays

Use this procedure to obtain values for the `vid` and `pid` attributes that are used with other forms of the `vxddladm` command.

To list all supported disk arrays

- ◆ Type the following command:

```
# vxddladm listsupport all
```

Displaying details about a supported array library

DMP enables you to display details about the Array Support Libraries (ASL).

The Array Support Libraries are in the directory `/etc/vx/lib/discovery.d`.

To display details about a supported array library

- ◆ Type the following command:

```
# vxddladm listsupport libname=library_name.so
```

This command displays the vendor ID (`VID`), product IDs (`PIDs`) for the arrays, array types (for example, `A/A` or `A/P`), and array names. The following is sample output.

```
# vxddladm listsupport libname=libvxfujitsu.so
ATTR_NAME          ATTR_VALUE
=====
LIBNAME             libvxfujitsu.so
VID                 vendor
PID                 GR710, GR720, GR730
                   GR740, GR820, GR840
ARRAY_TYPE          A/A, A/P
ARRAY_NAME          FJ_GR710, FJ_GR720, FJ_GR730
                   FJ_GR740, FJ_GR820, FJ_GR840
```

Excluding support for a disk array library

You can exclude support for disk arrays that depends on a particular disk array library. You can also exclude support for disk arrays from a particular vendor.

To exclude support for a disk array library

- ◆ To exclude support for a disk array library, specify the array library to the following command.

```
# vxddladm excludearray libname=libvxemc.so
```

You can also exclude support for disk arrays from a particular vendor, as shown in this example:

```
# vxddladm excludearray vid=ACME pid=X1
```

Re-including support for an excluded disk array library

If you previously excluded support for all arrays that depend on a particular disk array library, use this procedure to include the support for those arrays. This procedure removes the library from the exclude list.

To re-include support for an excluded disk array library

- ◆ If you have excluded support for all arrays that depend on a particular disk array library, you can use the `includearray` keyword to remove the entry from the exclude list, as shown in the following example:

```
# vxddladm includearray libname=libvxemc.so
```

This command adds the array library to the database so that the library can once again be used in device discovery. If `vxconfigd` is running, you can use the `vxdisk scandisks` command to discover the arrays and add their details to the database.

Listing excluded disk arrays

To list all disk arrays that are currently excluded from use by VxVM

- ◆ Type the following command:

```
# vxddladm listexclude
```

Listing supported disks in the DISKS category

To list disks that are supported in the `DISKS` (JBOD) category

- ◆ Type the following command:

```
# vxddladm listjbod
```

Adding unsupported disk arrays to the DISKS category

Disk arrays should be added as JBOD devices if no Array Support Library (ASL) is available for the array.

JBODs are assumed to be Active/Active (A/A) unless otherwise specified. If a suitable ASL is not available, an A/A-A, A/P or A/PF array must be claimed as an Active/Passive (A/P) JBOD to prevent path delays and I/O failures. If a JBOD is ALUA-compliant, it is added as an ALUA array.

See [“How DMP works”](#) on page 170.

Warning: The procedure in this section ensures that Dynamic Multi-Pathing (DMP) is set up correctly on an array that is not supported by Veritas Volume Manager.

If DMP is not set up correctly, Veritas Volume Manager treats the independent paths to the disks as separate devices, which can result in data corruption.

To add an unsupported disk array to the DISKS category

- 1 Use the following command to identify the vendor ID and product ID of the disks in the array:

```
# /etc/vx/diag.d/vxscsiinq device_name
```

where *device_name* is the device name of one of the disks in the array. Note the values of the vendor ID (VID) and product ID (PID) in the output from this command. For Fujitsu disks, also note the number of characters in the serial number that is displayed.

The following example output shows that the vendor ID is SEAGATE and the product ID is ST318404LSUN18G.

```
Vendor id (VID)      : SEAGATE
Product id (PID)     : ST318404LSUN18G
Revision            : 8507
Serial Number       : 0025T0LA3H
```

- 2 Stop all applications, such as databases, from accessing VxVM volumes that are configured on the array, and unmount all file systems and Storage Checkpoints that are configured on the array.
- 3 If the array is of type A/A-A, A/P or A/PF, configure it in autotrespass mode.

- 4 Enter the following command to add a new JBOD category:

```
# vxddladm addjbod vid=vendorid [pid=productid] \  
[serialnum=opcode/pagecode/offset/length] \  
[cabinetnum=opcode/pagecode/offset/length] policy={aa|ap}]
```

where *vendorid* and *productid* are the VID and PID values that you found from the previous step. For example, *vendorid* might be FUJITSU, IBM, or SEAGATE. For Fujitsu devices, you must also specify the number of characters in the serial number as the argument to the *length* argument (for example, 10). If the array is of type A/A-A, A/P or A/PF, you must also specify the *policy=ap* attribute.

Continuing the previous example, the command to define an array of disks of this type as a JBOD would be:

```
# vxddladm addjbod vid=SEAGATE pid=ST318404LSUN18G
```

- 5 Use the `vxctl enable` command to bring the array under VxVM control.

```
# vxctl enable
```

See [“Enabling discovery of new disk arrays”](#) on page 247.

- 6 To verify that the array is now supported, enter the following command:

```
# vxddladm listjbod
```

The following is sample output from this command for the example array:

VID	PID	SerialNum (Cmd/PageCode/off/len)	CabinetNum (Cmd/PageCode/off/len)	Policy
SEAGATE	ALL PIDs	18/-1/36/12	18/-1/10/11	Disk
SUN	SESS01	18/-1/36/12	18/-1/12/11	Disk

- 7 To verify that the array is recognized, use the `vxddmpadm listenclosure` command as shown in the following sample output for the example array:

```
# vxddmpadm listenclosure
ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS    ARRAY_TYPE LUN_COUNT
=====
Disk        Disk        DISKS    CONNECTED Disk        2
```

The enclosure name and type for the array are both shown as being set to `Disk`. You can use the `vxddisk list` command to display the disks in the array:

```
# vxddisk list
DEVICE      TYPE        DISK      GROUP      STATUS
Disk_0      auto:none   -         -          online invalid
Disk_1      auto:none   -         -          online invalid
...
```

- 8 To verify that the DMP paths are recognized, use the `vxddmpadm getdmpnode` command as shown in the following sample output for the example array:

```
# vxddmpadm getdmpnode enclosure=Disk
NAME      STATE    ENCLR-TYPE PATHS ENBL DSBL ENCLR-NAME
=====
Disk_0    ENABLED  Disk        2    2    0    Disk
Disk_1    ENABLED  Disk        2    2    0    Disk
...
```

The output in this example shows that there are two paths to the disks in the array.

For more information, enter the command `vxddladm help addjbod`.

See the `vxddladm(1M)` manual page.

See the `vxddmpadm(1M)` manual page.

Removing disks from the DISKS category

Use the procedure in this section to remove disks from the DISKS category.

To remove disks from the DISKS category

- ◆ Use the `vxddladm` command with the `rmjbod` keyword. The following example illustrates the command for removing disks which have the vendor id of SEAGATE:

```
# vxddladm rmjbod vid=SEAGATE
```


Foreign devices

DDL may not be able to discover some devices that are controlled by third-party drivers, such as those that provide multi-pathing or RAM disk capabilities. For these devices it may be preferable to use the multi-pathing capability that is provided by the third-party drivers for some arrays rather than using Dynamic Multi-Pathing (DMP). Such foreign devices can be made available as simple disks to VxVM by using the `vxddladm addforeign` command. This also has the effect of bypassing DMP for handling I/O. The following example shows how to add entries for block and character devices in the specified directories:

```
# vxddladm addforeign blockdir=/dev/foo/dsk \  
    chardir=/dev/foo/rdisk
```

By default, this command suppresses any entries for matching devices in the OS-maintained device tree that are found by the autodiscovery mechanism. You can override this behavior by using the `-f` and `-n` options as described on the `vxddladm(1M)` manual page.

After adding entries for the foreign devices, run the command to discover the devices as simple disks. These disks then behave in the same way as autoconfigured disks.

Use either the `vxdisk scandisks` or the `vxctl enable` command to discover the devices.

The foreign device feature was introduced in VxVM 4.0 to support non-standard devices such as RAM disks, some solid state disks, and pseudo-devices such as EMC PowerPath.

Foreign device support has the following limitations:

- A foreign device is always considered as a disk with a single path. Unlike an autodiscovered disk, it does not have a DMP node.
- It is not supported for shared disk groups in a clustered environment. Only standalone host systems are supported.
- It is not supported for Persistent Group Reservation (PGR) operations.
- It is not under the control of DMP, so enabling of a failed disk cannot be automatic, and DMP administrative commands are not applicable.
- Enclosure information is not available to VxVM. This can reduce the availability of any disk groups that are created using such devices.
- The I/O Fencing and Cluster File System features are not supported for foreign devices.

If a suitable ASL is available and installed for an array, these limitations are removed.

See [“Third-party driver coexistence”](#) on page 247.

Making devices invisible to VxVM

Use this procedure to exclude a device from the view of VxVM. The options to prevent a device from being multi-pathed by the VxVM DMP driver (`vxvmp`) are deprecated.

To make devices invisible to VxVM

- 1 Run the `vxdiskadm` command, and select `Prevent multipathing/Suppress devices from VxVM's view` from the main menu. You are prompted to confirm whether you want to continue.
- 2 Select the operation you want to perform from the following options:

Option 1	Suppresses all paths through the specified controller from the view of VxVM.
Option 2	Suppresses specified paths from the view of VxVM.
Option 3	Suppresses disks from the view of VxVM that match a specified Vendor ID and Product ID combination. The root disk cannot be suppressed. The operation fails if the VID:PID of an external disk is the same VID:PID as the root disk and the root disk is encapsulated under VxVM.
Option 4	Suppresses all but one path to a disk. Only one path is made visible to VxVM. This operation is deprecated, since it can lead to unsupported configurations.
Option 5	Prevents multi-pathing for all disks on a specified controller by VxVM. This operation is deprecated, since it can lead to unsupported configurations.
Option 6	Prevents multi-pathing of a disk by VxVM. The disks that correspond to a specified path are claimed in the OTHER_DISKS category and are not multi-pathed. This operation is deprecated, since it can lead to unsupported configurations.
Option 7	Prevents multi-pathing for the disks that match a specified Vendor ID and Product ID combination. The disks that correspond to a specified Vendor ID and Product ID combination are claimed in the OTHER_DISKS category and are not multi-pathed. This operation is deprecated, since it can lead to unsupported configurations.
Option 8	Lists the devices that are currently suppressed.

Making devices visible to VxVM

Use this procedure to make a device visible to VxVM again. The options to allow multi-pathing by the VxVM DMP driver (`vxddmp`) are deprecated.

To make devices visible to VxVM

- 1 Run the `vxdiskadm` command, and select `Allow multipathing/Unsuppress devices from VxVM's view` from the main menu. You are prompted to confirm whether you want to continue.
- 2 Select the operation you want to perform from the following options:

Option 1	Unsuppresses all paths through the specified controller from the view of VxVM.
Option 2	Unsuppresses specified paths from the view of VxVM.
Option 3	Unsuppresses disks from the view of VxVM that match a specified Vendor ID and Product ID combination.
Option 4	Removes a pathgroup definition. (A pathgroup explicitly defines alternate paths to the same disk.) Once a pathgroup has been removed, all paths that were defined in that pathgroup become visible again. This operation is deprecated.
Option 5	Allows multi-pathing of all disks that have paths through the specified controller. This operation is deprecated.
Option 6	Allows multi-pathing of a disk by VxVM. This operation is deprecated.
Option 7	Allows multi-pathing of disks that match a specified Vendor ID and Product ID combination. This operation is deprecated.
Option 8	Lists the devices that are currently suppressed.

About enabling and disabling I/O for controllers and storage processors

DMP lets you to turn off I/O through an HBA controller or the array port of a storage processor so that you can perform administrative operations. This feature can be

used for maintenance of HBA controllers on the host, or array ports that are attached to disk arrays supported by SFCFSA. I/O operations to the HBA controller or the array port can be turned back on after the maintenance task is completed. You can accomplish these operations using the `vxddmpadm` command.

For Active/Active type disk arrays, when you disable the I/O through an HBA controller or array port, the I/O continues on the remaining paths. For Active/Passive type disk arrays, if disabling I/O through an HBA controller or array port resulted in all primary paths being disabled, DMP will failover to secondary paths and I/O will continue on them.

DMP does not support the operations to enable I/O or disable I/O for the controllers that use Third-Party Drivers (TPD) for multi-pathing.

After the administrative operation is over, use the `vxddmpadm` command to re-enable the paths through the HBA controllers.

See [“Disabling I/O for paths, controllers or array ports”](#) on page 298.

See [“Enabling I/O for paths, controllers or array ports”](#) on page 299.

Note: From release 5.0 of VxVM, these operations are supported for controllers that are used to access disk arrays on which cluster-shareable disk groups are configured.

You can also perform certain reconfiguration operations dynamically online.

See [“About online dynamic reconfiguration”](#) on page 310.

About displaying Veritas Dynamic Multi-Pathing (DMP) information

You can use the `vxddmpadm` command to list DMP database information and perform other administrative tasks. This command allows you to list all controllers that are connected to disks, and other related information that is stored in the DMP database. You can use this information to locate system hardware, and to help you decide which controllers need to be enabled or disabled.

The `vxddmpadm` command also provides useful information such as disk array serial numbers, which DMP devices (disks) are connected to the disk array, and which paths are connected to a particular controller, enclosure or array port.

You can also use the `vxddisk` command to display the multi-pathing information for a particular metadvice. The metadvice is a device representation of a physical disk having multiple physical paths through the system's HBA controllers. In DMP,

all the physical disks in the system are represented as metadevices with one or more physical paths.

Displaying Veritas Dynamic Multi-Pathing (DMP) information includes the following topics:

See [“Retrieving information about a DMP node”](#) on page 271.

See [“Displaying consolidated information about the DMP nodes”](#) on page 272.

See [“Displaying the members of a LUN group”](#) on page 273.

See [“Displaying paths controlled by a DMP node, controller, enclosure, or array port”](#) on page 274.

See [“Displaying information about controllers”](#) on page 276.

See [“Displaying information about enclosures”](#) on page 277.

See [“Displaying extended device attributes”](#) on page 280.

See [“Displaying the paths to a disk”](#) on page 266.

See [“Administering DMP using vxdkmpadm”](#) on page 269.

Displaying the paths to a disk

See [“About displaying Veritas Dynamic Multi-Pathing \(DMP\) information”](#) on page 265.

The `vxdisk` command is used to display the multi-pathing information for a particular metadvice. The metadvice is a device representation of a physical disk having multiple physical paths through the system’s HBA controllers. In DMP, all the physical disks in the system are represented as metadevices with one or more physical paths.

To display the multi-pathing information on a system

- ◆ Use the `vxdisk path` command to display the relationships between the device paths, disk access names, disk media names and disk groups on a system as shown here:

```
# vxdisk path
```

SUBPATH	DANAME	DMNAME	GROUP	STATE
c1t0d0s2	c1t0d0s2	mydg01	mydg	ENABLED
c4t0d0s2	c1t0d0s2	mydg01	mydg	ENABLED
c1t1d0s2	c1t1d0s2	mydg02	mydg	ENABLED
c4t1d0s2	c1t1d0s2	mydg02	mydg	ENABLED
.	.	.	.	.

This shows that two paths exist to each of the two disks, `mydg01` and `mydg02`, and also indicates that each disk is in the `ENABLED` state.

To view multi-pathing information for a particular metadvice

- 1 Use the following command to view multi-pathing information for the specified device:

```
# vxdisk list devicename
```

For example, the command output displays multi-pathing information for the device `c2t0d0s2` as follows:

```
# vxdisk list c2t0d0s2

Device      c2t0d0
devicetag   c2t0d0
type        sliced
hostid      system01
.
.
.
Multipathing information:
numpaths:   2
c2t0d0s2    state=enabled   type=primary
c1t0d0s2    state=disabled  type=secondary
```

In the sample output above, the `numpaths` line shows that there are 2 paths to the device. The next two lines in the "Multipathing information" section show that one path is active (`state=enabled`) and that the other path has failed (`state=disabled`).

The `type` field is shown for disks on Active/Passive type disk arrays such as the EMC CLARiiON, Hitachi HDS 9200 and 9500, Sun StorEdge 6xxx, and Sun StorEdge T3 array. This field indicates the primary and secondary paths to the disk.

The `type` field is not displayed for disks on Active/Active type disk arrays such as the EMC Symmetrix, Hitachi HDS 99xx and Sun StorEdge 99xx Series, and IBM ESS Series. Such arrays have no concept of primary and secondary paths.

- 2 Alternately, you can use the following command to view multi-pathing information:

```
# vxddmpadm getsubpaths dmpnodename=devicename
```

For example, to view multi-pathing information for `eva4k6k0_6`, use the following command:

```
# vxddmpadm getsubpaths dmpnodename=eva4k6k0_6
```

NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c0t50001FE1500A8F08d7s2	ENABLED (A)	PRIMARY	c0	EVA4K6K	eva4k6k0	-
c0t50001FE1500A8F09d7s2	ENABLED (A)	PRIMARY	c0	EVA4K6K	eva4k6k0	-
c0t50001FE1500A8F0Cd7s2	ENABLED	SECONDARY	c0	EVA4K6K	eva4k6k0	-
c0t50001FE1500A8F0Dd7s2	ENABLED	SECONDARY	c0	EVA4K6K	eva4k6k0	-

Administering DMP using vxddmpadm

The `vxddmpadm` utility is a command line administrative interface to DMP.

You can use the `vxddmpadm` utility to perform the following tasks:

- Retrieve the name of the DMP device corresponding to a particular path.
See [“Retrieving information about a DMP node”](#) on page 271.
- Display consolidated information about the DMP nodes
See [“Displaying consolidated information about the DMP nodes”](#) on page 272.
- Display the members of a LUN group.
See [“Displaying the members of a LUN group”](#) on page 273.
- List all paths under a DMP device node, HBA controller, enclosure, or array port.
See [“Displaying paths controlled by a DMP node, controller, enclosure, or array port”](#) on page 274.
- Display information about the HBA controllers on the host.
See [“Displaying information about controllers”](#) on page 276.
- Display information about enclosures.
See [“Displaying information about enclosures”](#) on page 277.
- Display information about array ports that are connected to the storage processors of enclosures.
See [“Displaying information about array ports”](#) on page 278.

- Display information about devices that are controlled by third-party multi-pathing drivers.
See [“Displaying DMP path information for devices under third-party driver control”](#) on page 279.
- Display extended devices attributes.
See [“Displaying extended device attributes”](#) on page 280.
- Suppress or include devices from VxVM control.
See [“Suppressing or including devices from VxVM control”](#) on page 282.
- Gather I/O statistics for a DMP node, enclosure, path or controller.
See [“Gathering and displaying DMP I/O statistics”](#) on page 283.
- Configure the attributes of the paths to an enclosure.
See [“Setting the attributes of the paths to an enclosure”](#) on page 289.
- Display the redundancy level of a device or enclosure
See [“Displaying the path redundancy level of a device or enclosure under DMP control”](#) on page 290.
- Specify the minimum number of active paths
See [“Specifying the minimum number of active paths for a device or enclosure under DMP control”](#) on page 290.
- Display or set the I/O policy that is used for the paths to an enclosure.
See [“Specifying the DMP I/O policy”](#) on page 292.
- Enable or disable I/O for a path, HBA controller or array port on the system.
See [“Disabling I/O for paths, controllers or array ports”](#) on page 298.
- Rename an enclosure.
See [“Renaming an enclosure”](#) on page 300.
- Configure how DMP responds to I/O request failures.
See [“Configuring DMP’s response to I/O failures”](#) on page 301.
- Configure the I/O throttling mechanism.
See [“Configuring the DMP I/O throttling mechanism”](#) on page 302.
- Control the operation of the DMP path restoration thread.
See [“Configuring DMP path restoration policies”](#) on page 305.
- Configure array policy modules
See [“Configuring array policy modules”](#) on page 307.
- Get or set the values of various tunables used by DMP.
See [“DMP tunable parameters”](#) on page 884.

The following sections cover these tasks in detail along with sample output.

See the `vxddmpadm(1M)` manual page.

Retrieving information about a DMP node

The following command displays the DMP node that controls a particular physical path:

```
# vxddmpadm getdmpnode nodename=pathname
```

The physical path is specified by argument to the `nodename` attribute, which must be a valid path listed in the device directory.

The device directory is the `/dev/rdisk` directory.

The command displays output similar to the following example output.

```
# vxddmpadm getdmpnode nodename=c0t5006016041E03B33d0s2
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
emc_clariion0_16	ENABLED	EMC_CLARIION	6	6	0	emc_clariion0

Use the `-v` option to display the LUN serial number and the array volume ID.

```
# vxddmpadm -v getdmpnode nodename=c0t5006016041E03B33d0s2
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME	SERIAL-NO	ARRAY_VOL_ID
emc_clariion0_16	ENABLED	EMC_CLARIION	6	6	0	emc_clariion0	600601606	16

Use the `enclosure` attribute with `getdmpnode` to obtain a list of all DMP nodes for the specified enclosure.

```
# vxddmpadm getdmpnode enclosure=enc0
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
c2t1d0s2	ENABLED	T300	2	2	0	enc0
c2t1d1s2	ENABLED	T300	2	2	0	enc0
c2t1d2s2	ENABLED	T300	2	2	0	enc0
c2t1d3s2	ENABLED	T300	2	2	0	enc0

Use the `dmpnodename` attribute with `getdmpnode` to display the DMP information for a given DMP node.

```
# vxddmpadm getdmpnode dmpnodename=emc_clariion0_158
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
=====						
emc_clariion0_158	ENABLED	EMC_CLARiion	1	1	0	emc_clariion0

Displaying consolidated information about the DMP nodes

The `vxddpdm list dmpnode` command displays the detail information of a DMP node. The information includes the enclosure name, LUN serial number, port id information, device attributes, etc.

The following command displays the consolidated information for all of the DMP nodes in the system:

```
# vxddpdm list dmpnode all
```

Use the `enclosure` attribute with `list dmpnode` to obtain a list of all DMP nodes for the specified enclosure.

```
# vxddpdm list dmpnode enclosure=enclosure name
```

For example, the following command displays the consolidated information for all of the DMP nodes in the `enc0` enclosure.

```
# vxddpdm list dmpnode enclosure=enc0
```

Use the `dmpnodename` attribute with `list dmpnode` to display the DMP information for a given DMP node. The DMP node can be specified by name or by specifying a path name. The detailed information for the specified DMP node includes path information for each subpath of the listed dmpnode.

The path state differentiates between a path that is disabled due to a failure and a path that has been manually disabled for administrative purposes. A path that has been manually disabled using the `vxddpdm disable` command is listed as `disabled(m)`.

```
# vxddpdm list dmpnode dmpnodename=dmpnodename
```

For example, the following command displays the consolidated information for the DMP node `emc_clariion0_158`.

```
# vxddpdm list dmpnode dmpnodename=emc_clariion0_158
```

```
dmpdev = emc_clariion0_158
state   = enabled
enclosure = emc_clariion0
cab-sno  = CK200070400359
asl      = libvxCLARiion.so
```

```

vid          = DGC
pid          = DISK
array-name   = EMC_CLARiiON
array-type   = CLR-A/PF
iopolicy     = MinimumQ
avid         = 158
lun-sno      = 600601606D121B008FB6E0CA8EDBDB11
udid         = DGC%5FDISK%5FCK200070400359%5F600601606D121B008FB6E0CA8EDBDB11
dev-attr     = lun
###path      = name state type transport ctlr hwpath aprotID aprotWWN attr
path         = c0t5006016141E03B33d1s2 enabled(a) primary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 A5 50:06:01:61:41:e0:3b:33 -
path         = c0t5006016041E03B33d1s2 enabled(a) primary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 A4 50:06:01:60:41:e0:3b:33 -
path         = c0t5006016841E03B33d1s2 enabled secondary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 B4 50:06:01:68:41:e0:3b:33 -
path         = c1t5006016141E03B33d1s2 enabled(a) primary FC c1
/pci@1e,600000/SUNW,emlxs@3,1/fp@0,0 A5 50:06:01:61:41:e0:3b:33 -
path         = c1t5006016841E03B33d1s2 enabled secondary FC c1
/pci@1e,600000/SUNW,emlxs@3,1/fp@0,0 B4 50:06:01:68:41:e0:3b:33 -
path         = c1t5006016041E03B33d1s2 enabled(a) primary FC c1
/pci@1e,600000/SUNW,emlxs@3,1/fp@0,0 A4 50:06:01:60:41:e0:3b:33 -

```

Displaying the members of a LUN group

The following command displays the DMP nodes that are in the same LUN group as a specified DMP node:

```
# vxddpadm getlungroup dmpnodename=dmpnodename
```

The above command displays output such as the following.

```
# vxddpadm getlungroup dmpnodename=c11t0d10s2
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
c11t0d8s2	ENABLED	ACME	2	2	0	enc1
c11t0d9s2	ENABLED	ACME	2	2	0	enc1
c11t0d10s2	ENABLED	ACME	2	2	0	enc1
c11t0d11s2	ENABLED	ACME	2	2	0	enc1

Displaying paths controlled by a DMP node, controller, enclosure, or array port

The `vxddmpadm getsubpaths` command lists all of the paths known to DMP. The `vxddmpadm getsubpaths` command also provides options to list the subpaths through a particular DMP node, controller, enclosure, or array port. To list the paths through an array port, specify either a combination of enclosure name and array port id, or array port WWN.

To list all subpaths known to DMP:

```
# vxddmpadm getsubpaths
```

NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-NAME	CTLR	ATTRS
=====						
c1t65d0s2	ENABLED (A)	-	Disk_1	Disk	c1	-
c1t66d0s2	ENABLED (A)	-	Disk_2	Disk	c1	-
c2t65d0s2	ENABLED (A)	-	Disk_1	Disk	c2	-
c2t66d0s2	ENABLED (A)	-	Disk_2	Disk	c2	-
c3t2d0s2	ENABLED (A)	-	EMC0_1	EMC0	c3	-
c3t2d1s2	ENABLED (A)	-	EMC0_2	EMC0	c3	-
c4t2d0s2	ENABLED (A)	-	EMC0_1	EMC0	c4	-
c4t2d1s2	ENABLED (A)	-	EMC0_2	EMC0	c4	-

The `vxddmpadm getsubpaths` command combined with the `dmpnodename` attribute displays all the paths to a LUN that are controlled by the specified DMP node name from the `/dev/vx/rddmp` directory:

```
# vxddmpadm getsubpaths dmpnodename=c2t66d0s2
```

NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
=====						
c2t66d0s2	ENABLED (A)	PRIMARY	c2	ACME	enc0	-
c1t66d0s2	ENABLED	PRIMARY	c1	ACME	enc0	-

For A/A arrays, all enabled paths that are available for I/O are shown as `ENABLED (A)`.

For A/P arrays in which the I/O policy is set to `singleactive`, only one path is shown as `ENABLED (A)`. The other paths are enabled but not available for I/O. If the I/O policy is not set to `singleactive`, DMP can use a group of paths (all primary or all secondary) for I/O, which are shown as `ENABLED (A)`.

See “Specifying the DMP I/O policy” on page 292.

Paths that are in the `DISABLED` state are not available for I/O operations.

A path that was manually disabled by the system administrator displays as DISABLED(M). A path that failed displays as DISABLED.

You can use `getsubpaths` to obtain information about all the paths that are connected to a particular HBA controller:

```
# vxddpdm getsubpaths ctrl=c2
```

NAME	STATE [-]	PATH-TYPE [-]	CTRL-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS
c2t1d0s2	ENABLED (A)	PRIMARY	c2t1d0s2	ACME	enc0	-
c2t2d0s2	ENABLED (A)	PRIMARY	c2t2d0s2	ACME	enc0	-
c2t3d0s2	DISABLED	SECONDARY	c2t3d0s2	ACME	enc0	-
c2t4d0s2	ENABLED	SECONDARY	c2t4d0s2	ACME	enc0	-

You can also use `getsubpaths` to obtain information about all the paths that are connected to a port on an array. The array port can be specified by the name of the enclosure and the array port ID, or by the worldwide name (WWN) identifier of the array port:

```
# vxddpdm getsubpaths enclosure=enclosure portid=portid
# vxddpdm getsubpaths pwwn=pwwn
```

For example, to list subpaths through an array port through the enclosure and the array port ID:

```
# vxddpdm getsubpaths enclosure=HDS9500-ALUA0 portid=1A
```

NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-NAME	CTRL	ATTRS
c1t65d0s2	ENABLED (A)	PRIMARY	c1t65d0s2	HDS9500-ALUA0	c1	-
c1t66d0s2	ENABLED (A)	PRIMARY	c1t66d0s2	HDS9500-ALUA0	c1	-

For example, to list subpaths through an array port through the WWN:

```
# vxddpdm getsubpaths pwwn=20:00:00:E0:8B:06:5F:19]
```

NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-NAME	CTRL	ATTRS
c1t65d0s2	ENABLED (A)	PRIMARY	c1t65d0s2	HDS9500-ALUA0	c1	-
c1t66d0s2	ENABLED (A)	PRIMARY	c1t66d0s2	HDS9500-ALUA0	c1	-

You can use `getsubpaths` to obtain information about all the subpaths of an enclosure.

```
# vxddpdm getsubpaths enclosure=enclosure_name [ctrl=ctrlname]
```

To list all subpaths of an enclosure:

```
# vxddmpadm getsubpaths enclosure=Disk
```

NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-NAME	CTLR	ATTRS
c1t65d0s2	ENABLED (A)	-	Disk_1	Disk	c1	-
c1t66d0s2	ENABLED (A)	-	Disk_2	Disk	c1	-
c2t65d0s2	ENABLED (A)	-	Disk_1	Disk	c2	-
c2t66d0s2	ENABLED (A)	-	Disk_2	Disk	c2	-

To list all subpaths of a controller on an enclosure:

```
# vxddmpadm getsubpaths enclosure=Disk ctlr=c1
```

By default, the output of the `vxddmpadm getsubpaths` command is sorted by enclosure name, DMP node name, and within that, path name.

To sort the output based on the pathname, the DMP node name, the enclosure name, or the host controller name, use the `-s` option.

To sort subpaths information, use the following command:

```
# vxddmpadm -s {path | dmpnode | enclosure | ctlr} getsubpaths \
[all | ctlr=ctlr_name | dmpnodename=dmp_device_name | \
enclosure=enclr_name [ctlr=ctlr_name | portid=array_port_ID] | \
pwwn=port_WWN | tpdnodename=tpd_node_name]
```

Displaying information about controllers

The following command lists attributes of all HBA controllers on the system:

```
# vxddmpadm listctlr all
```

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
c1	OTHER	ENABLED	other0
c2	X1	ENABLED	jbod0
c3	ACME	ENABLED	enc0
c4	ACME	ENABLED	enc0

This output shows that the controller `c1` is connected to disks that are not in any recognized DMP category as the enclosure type is `OTHER`.

The other controllers are connected to disks that are in recognized DMP categories.

All the 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 using the `vxddmpadm disable` command.

The following forms of the command lists controllers belonging to a specified enclosure or enclosure type:

```
# vxddmpadm listctlr enclosure=enc0
```

or

```
# vxddmpadm listctlr type=ACME
```

CTLR-NAME	ENCLR-TYPE	STATE	ENCLR-NAME
=====			
c2	ACME	ENABLED	enc0
c3	ACME	ENABLED	enc0

The `vxddmpadm getctlr` command displays HBA vendor details and the Controller ID. For iSCSI devices, the Controller ID is the IQN or IEEE-format based name. For FC devices, the Controller ID is the WWN. Because the WWN is obtained from ESD, this field is blank if ESD is not running. ESD is a daemon process used to notify DDL about occurrence of events. The WWN shown as 'Controller ID' maps to the WWN of the HBA port associated with the host controller.

```
# vxddmpadm getctlr c5
```

LNAME	PNAME	VENDOR	CTLR-ID
=====			
c5	c5	qlogic	20:07:00:a0:b8:17:e1:37

Displaying information about enclosures

To display the attributes of a specified enclosure, including its enclosure type, enclosure serial number, status, array type, and number of LUNs, use the following command:

```
# vxddmpadm listenclosure enc0
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE	LUN_COUNT
=====					
enc0	T3	60020f20000001a90000	CONNECTED	A/P	30

The following command lists attributes for all enclosures in a system:

```
# vxddmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE	LUN_COUNT
Disk	Disk	DISKS	CONNECTED	Disk	6
SENA0	SENA	508002000001d660	CONNECTED	A/A	57
enc0	T3	60020f20000001a90000	CONNECTED	A/P	30

If an A/P or ALUA array is under the control of MPxIO, then DMP claims the devices in A/A mode. The output of the above commands shows the ARRAY_TYPE as A/A. For arrays under MPxIO control, DMP does not store A/P-specific attributes or ALUA-specific attributes. These attributes include primary/secondary paths, port serial number, and the array controller ID.

Displaying information about array ports

Use the commands in this section to display information about array ports. The information displayed for an array port includes the name of its enclosure, and its ID and worldwide name (WWN) identifier.

Note: DMP does not report information about array ports for LUNs that are controlled by the native multi-pathing driver. DMP reports pWWN information only if the dmp_monitor_fabric tunable is on, and the event source daemon (esd) is running.

To display the attributes of an array port that is accessible via a path, DMP node or HBA controller, use one of the following commands:

```
# vxddmpadm getportids path=path-name
# vxddmpadm getportids dmpnodename=dmpnode-name
# vxddmpadm getportids ctlr=ctlr-name
```

The following form of the command displays information about all of the array ports within the specified enclosure:

```
# vxddmpadm getportids enclosure=enclr-name
```

The following example shows information about the array port that is accessible via DMP node c2t66d0s2:

```
# vxddmpadm getportids dmpnodename=c2t66d0s2
```

NAME	ENCLR-NAME	ARRAY-PORT-ID	pWWN
c2t66d0s2	HDS9500V0	1A	20:00:00:E0:8B:06:5F:19

Displaying DMP path information for devices under third-party driver control

The third-party driver (TPD) coexistence feature allows I/O that is controlled by third-party multi-pathing drivers to bypass DMP while retaining the monitoring capabilities of DMP. The following commands allow you to display the paths that DMP has discovered for a given TPD device, and the TPD device that corresponds to a given TPD-controlled node discovered by DMP:

```
# vxddmpadm getsubpaths tpdnodename=TPD_node_name
# vxddmpadm gettpdnode nodename=TPD_path_name
```

See [“Changing device naming for TPD-controlled enclosures”](#) on page 330.

For example, consider the following disks in an EMC Symmetrix array controlled by PowerPath, which are known to DMP:

```
# vxddisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
emcpower10s2	auto:sliced	disk1	ppdg	online
emcpower11s2	auto:sliced	disk2	ppdg	online
emcpower12s2	auto:sliced	disk3	ppdg	online
emcpower13s2	auto:sliced	disk4	ppdg	online
emcpower14s2	auto:sliced	disk5	ppdg	online
emcpower15s2	auto:sliced	disk6	ppdg	online
emcpower16s2	auto:sliced	disk7	ppdg	online
emcpower17s2	auto:sliced	disk8	ppdg	online
emcpower18s2	auto:sliced	disk9	ppdg	online
emcpower19s2	auto:sliced	disk10	ppdg	online

The following command displays the paths that DMP has discovered, and which correspond to the PowerPath-controlled node, emcpower10s2:

```
# vxddmpadm getsubpaths tpdnodename=emcpower10s2
```

NAME	TPDNODENAME	PATH-TYPE [-]	DMP-NODENAME	ENCLR-TYPE	ENCLR-NAME
c7t0d10s2emcpower10s2-			emcpower10s2	EMC	EMC0
c6t0d10s2emcpower10s2-			emcpower10s2	EMC	EMC0

Conversely, the next command displays information about the PowerPath node that corresponds to the path, c7t0d10s2, discovered by DMP:

```
# vxddmpadm gettpdnode nodename=c7t0d10s2
```

NAME	STATE	PATHS	ENCLR-TYPE	ENCLR-NAME
=====				
emcpower10s2	ENABLED	2	EMC	EMC0

Displaying extended device attributes

Device Discovery Layer (DDL) extended attributes are attributes or flags corresponding to a VxVM or DMP LUN or Disk and which are discovered by DDL. These attributes identify a LUN to a specific hardware category.

The list of categories includes:

Hardware RAID types	Displays what kind of Storage RAID Group the LUN belongs to
Thin Provisioning Discovery and Reclamation	Displays the LUN's thin reclamation abilities
Device Media Type	Displays the type of media –whether SSD (solid state disk)
Storage-based Snapshot/Clone	Displays whether the LUN is a SNAPSHOT or a CLONE of a PRIMARY LUN
Storage-based replication	Displays if the LUN is part of a replicated group across a remote site
Transport	Displays what kind of HBA is used to connect to this LUN (FC, SATA, iSCSI)

Each LUN can have one or more of these extended attributes. DDL discovers the extended attributes during device discovery from the array support library (ASL). If Veritas Operations Manager (VOM) is present, DDL can also obtain extended attributes from the VOM Management Server for hosts that are configured as managed hosts.

The `vxdisk -p list` command displays DDL extended attributes. For example, the following command shows attributes of “std”, “fc”, and “RAID_5” for this LUN:

```
# vxdisk -p list
DISK          : tagmastore-usp0_0e18
DISKID        : 1253585985.692.rx2600h11
VID           : HITACHI
UDID          : HITACHI%5FOPEN-V%5F02742%5F0E18
REVISION      : 5001
PID           : OPEN-V
PHYS_CTLR_NAME : 0/4/1/1.0x50060e8005274246
```

```

LUN_SNO_ORDER   : 411
LUN_SERIAL_NO   : 0E18
LIBNAME         : libvxhdsusp.sl
HARDWARE_MIRROR : no
DMP_DEVICE      : tagmastore-usp0_0e18
DDL_THIN_DISK   : thick
DDL_DEVICE_ATTR : std fc RAID_5
CAB_SERIAL_NO   : 02742
ATYPE          : A/A
ARRAY_VOLUME_ID : 0E18
ARRAY_PORT_PWWN : 50:06:0e:80:05:27:42:46
ANAME          : TagmaStore-USP
TRANSPORT      : FC

```

The `vxdisk -x attribute -p list` command displays the one-line listing for the property list and the attributes. The following example shows two Hitachi LUNs that support Thin Reclamation via the attribute `hdprclm`:

```

# vxdisk -x DDL_DEVICE_ATTR -p list
DEVICE                DDL_DEVICE_ATTR
tagmastore-usp0_0a7a  std fc RAID_5
tagmastore-usp0_065a  hdprclm fc
tagmastore-usp0_065b  hdprclm fc

```

User can specify multiple `-x` options in the same command to display multiple entries. For example:

```

# vxdisk -x DDL_DEVICE_ATTR -x VID -p list
DEVICE                VID                DDL_DEVICE_ATTR
tagmastore-usp0_0a7a  HITACHI                std fc RAID_5
tagmastore-usp0_0a7b  HITACHI                std fc RAID_5
tagmastore-usp0_0a78  HITACHI                std fc RAID_5
tagmastore-usp0_0a79  HITACHI                std fc RAID_5
tagmastore-usp0_065a  HITACHI                hdprclm fc
tagmastore-usp0_065b  HITACHI                hdprclm fc
tagmastore-usp0_065c  HITACHI                hdprclm fc
tagmastore-usp0_065d  HITACHI                hdprclm fc

```

Use the `vxdisk -e list` command to show the `DDL_DEVICE_ATTR` property in the last column named `ATTR`.

```

# vxdisk -e list
DEVICE                TYPE  DISK  GROUP  STATUS  OS_NATIVE_NAME  ATTR
tagmastore-usp0_0a7a  auto  -    -      online  cl0t0d2         std fc RAID_5

```

```
tagmastore-usp0_0a7b auto - - online c10t0d3 std fc RAID_5
tagmastore-usp0_0a78 auto - - online c10t0d0 std fc RAID_5
tagmastore-usp0_0655 auto - - online c13t2d7 hdprclm fc
tagmastore-usp0_0656 auto - - online c13t3d0 hdprclm fc
tagmastore-usp0_0657 auto - - online c13t3d1 hdprclm fc
```

For a list of ASLs that supports Extended Attributes, and descriptions of these attributes, refer to the hardware compatibility list (HCL) at the following URL:

<http://www.symantec.com/docs/TECH170013>

Note: DMP does not support Extended Attributes for LUNs that are controlled by the native multi-pathing driver.

Suppressing or including devices from VxVM control

The `vxddmpadm exclude` command suppresses devices from VxVM based on the criteria that you specify. When a device is suppressed, DMP does not claim the device so that the device is not available for VxVM to use. You can add the devices back into VxVM control with the `vxddmpadm include` command. The devices can be included or excluded based on VID:PID combination, paths, controllers, or disks. You can use the bang symbol (!) to exclude or include any paths or controllers except the one specified.

The root disk cannot be suppressed.

The operation fails if the VID:PID of an external disk is the same VID:PID as the root disk and the root disk is encapsulated under VxVM.

Note: The ! character is a special character in some shells. The following syntax shows how to escape it in a bash shell.

```
# vxddmpadm exclude { all | product=VID:PID |
ctrlr=[\!]ctrlr | dmpnodename=diskname [ path=\!pathname] }

# vxddmpadm include { all | product=VID:PID |
ctrlr=[\!]ctrlr | dmpnodename=diskname [ path=\!pathname] }
```

where:

`all` – all devices

`product=VID:PID` – all devices with the specified VID:PID

`ctrlr=ctrlr` – all devices through the given controller

dmpnodename=*diskname* - all paths under the DMP node

dmpnodename=*diskname* path=*!pathname* - all paths under the DMP node except the one specified.

Gathering and displaying DMP I/O statistics

You can use the `vxddmpadm iostat` command to gather and display I/O statistics for a specified DMP node, enclosure, path or controller.

To enable the gathering of statistics, enter this command:

```
# vxddmpadm iostat start [memory=size]
```

To reset the I/O counters to zero, use this command:

```
# vxddmpadm iostat reset
```

The `memory` attribute can be used to limit the maximum amount of memory that is used to record I/O statistics for each CPU. The default limit is `32k` (32 kilobytes) per CPU.

To display the accumulated statistics at regular intervals, use the following command:

```
# vxddmpadm iostat show {all | ctrl=ctrl-name \  
    | dmpnodename=dmp-node \  
    | enclosure=enclr-name [portid=array-portid] \  
    | pathname=path-name | pwwn=array-port-wwn } \  
    [interval=seconds [count=N]]
```

This command displays I/O statistics for all paths (`all`), or for a specified controller, DMP node, enclosure, path or port ID. The statistics displayed are the CPU usage and amount of memory per CPU used to accumulate statistics, the number of read and write operations, the number of kilobytes read and written, and the average time in milliseconds per kilobyte that is read or written.

The `interval` and `count` attributes may be used to specify the interval in seconds between displaying the I/O statistics, and the number of lines to be displayed. The actual interval may be smaller than the value specified if insufficient memory is available to record the statistics.

To disable the gathering of statistics, enter this command:

```
# vxddmpadm iostat stop
```

Displaying cumulative DMP //O statistics

Use the `groupby` clause of the `vxddmpadm iostat` command to display cumulative I/O statistics listings per DMP node, controller, array port id, or host-array controller pair and enclosure. If the `groupby` clause is not specified, then the statistics are displayed per path.

By default, the read/write times are displayed in milliseconds up to 2 decimal places. The throughput data is displayed in terms of BLOCKS, and the output is scaled, meaning that the small values are displayed in small units and the larger values are displayed in bigger units, keeping significant digits constant. You can specify the units in which the statistics data is displayed. The `-u` option accepts the following options:

<code>h</code> or <code>H</code>	Displays throughput in the highest possible unit.
<code>k</code>	Displays throughput in kilobytes.
<code>m</code>	Displays throughput in megabytes.
<code>g</code>	Displays throughput in gigabytes.
<code>bytes</code> <code>b</code>	Displays throughput in exact number of bytes.
<code>us</code>	Displays average read/write time in microseconds.

To group by DMP node:

```
# vxddmpadm [-u unit] iostat show groupby=dmpnode \
[all | dmpnodename=dmpnodename | enclosure=enclr-name]
```

The following example shows grouping by DMP node.

```
# vxddmpadm iostat show groupby=dmpnode dmpnodename=c5t0d1s2
```

OPERATIONS	BLOCKS		AVG TIME (ms)			
DMPNODENAME	READS	WRITES	READS	WRITES	READS	WRITES
c5t0d1s2	0	0	0	0	0.00	0.00

To group by controller:

```
# vxddmpadm [-u unit] iostat show groupby=ctlr [ all | ctlr=ctlr ]
```

For example:

```
# vxddmpadm iostat show groupby=ctlr ctlr=c5
```


CTLRNAME	OPERATIONS		BLOCKS		AVG TIME (ms)	
	READS	WRITES	READS	WRITES	READS	WRITES
c5	224	14	54	7	4.20	11.10

To group by arrayport:

```
# vxddmpadm [-u unit] iostat show groupby=arrayport [ all \
| pwwn=array_pwwn | enclosure=enclr portid=array-port-id ]
```

For example:

```
# vxddmpadm -u m iostat show groupby=arrayport \
enclosure=HDS9500-ALUA0 portid=1A
```

PORTNAME	OPERATIONS		BYTES		AVG TIME (ms)	
	READS	WRITES	READS	WRITES	READS	WRITES
1A	743	1538	11m	24m	17.13	8.61

To group by enclosure:

```
# vxddmpadm [-u unit] iostat show groupby=enclosure [ all \
| enclosure=enclr ]
```

For example:

```
# vxddmpadm -u h iostat show groupby=enclosure enclosure=EMC_CLARiiON0
```

ENCLRNAME	OPERATIONS		BLOCKS		AVG TIME (ms)	
	READS	WRITES	READS	WRITES	READS	WRITES
EMC_CLARiiON	743	1538	11392k	24176k	17.13	8.61

You can also filter out entities for which all data entries are zero. This option is especially useful in a cluster environment which contains many failover devices. You can display only the statistics for the active paths.

To filter all zero entries from the output of the `iostat show` command:

```
# vxddmpadm [-u unit] -z iostat show [all|ctlr=ctlr_name |
dmpnodename=dmp_device_name | enclosure=enclr_name [portid=portid] |
pathname=path_name|pwwn=port_WWN] [interval=seconds [count=N]]
```

For example, to display the statistics for a dmpnodename and exclude the zero entries:

```
# vxddmpadm -z iostat show dmpnodename=c2t16d4s2
```

PATHNAME	OPERATIONS		BLOCKS		AVG TIME (ms)	
	READS	WRITES	READS	WRITES	READS	WRITES

```

c3t16d4s2  10      110      2      25      12.00      27.96
c2t17d4s2  20      126      4      29      9.50      19.41

```

To display average read/write times in microseconds.

```
# vxddmpadm -u us iostat show pathname=c2t17d4s2
```

PATHNAME	OPERATIONS		BLOCKS		AVG TIME(microsec)	
	READS	WRITES	READS	WRITES	READS	WRITES
c2t17d4s2	20	126	4	29	9500.00	19413.79

Displaying DMP statistics for queued or erroneous I/Os

Use the `vxddmpadm iostat show` command with the `-q` option to display the I/Os queued in DMP for a specified DMP node, or for a specified path or controller. For a DMP node, the `-q` option displays the I/Os on the specified DMP node that were sent to underlying layers. If a path or controller is specified, the `-q` option displays I/Os that were sent to the given path or controller and not yet returned to DMP.

See the `vxddmpadm(1m)` manual page for more information about the `vxddmpadm iostat` command.

To display queued I/O counts on a DMP node:

```
# vxddmpadm -q iostat show [filter] [interval=n [count=m]]
```

For example, to show the queued I/O counts on a DMP node:

```
# vxddmpadm -q iostat show dmpnodename=c5t2d1s2
```

DMPNODENAME	QUEUED I/Os		PENDING I/Os
	READS	WRITES	
c5t2d1s2	2	15	30

To display the count of I/Os that returned with errors on a DMP node, path or controller:

```
# vxddmpadm -e iostat show [filter] [interval=n [count=m]]
```

For example, to show the I/O counts that returned errors on a path:

```

cpu usage = 168144us      per cpu memory = 409600b
                           ERROR I/Os
PATHNAME      READS      WRITES
c1t2d1        0          0

```

```
# vxddmpadm -e iostat show pathname=c1t5006016041E03B33d6s2 interval=1
```

```

cpu usage = 168144us      per cpu memory = 409600b
                           ERROR I/Os
PATHNAME                READS    WRITES
c1t5006016041E03B33d6s2  0        0
c1t5006016041E03B33d6s2  0        0

```

Examples of gathering and displaying DMP I/O statistics using the vxddmpadm iostat command

The following is an example session using the vxddmpadm iostat command. The first command enables the gathering of I/O statistics:

```
# vxddmpadm iostat start
```

The next command displays the current statistics including the accumulated total numbers of read and write operations, and the kilobytes read and written, on all paths.

```

# vxddmpadm -u k iostat show all

                           cpu usage = 7952us      per cpu memory = 8192b
                           OPERATIONS              BYTES              AVG TIME (ms)
PATHNAME  READS    WRITES      READS      WRITES      READS    WRITES
c0t0d0    1088      0        557056k      0        0.00    0.00
c2t118d0   87       0        44544k      0        0.00    0.00
c3t118d0   0        0         0         0        0.00    0.00
c2t122d0   87       0        44544k      0        0.00    0.00
c3t122d0   0        0         0         0        0.00    0.00
c2t115d0   87       0        44544k      0        0.00    0.00
c3t115d0   0        0         0         0        0.00    0.00
c2t103d0   87       0        44544k      0        0.00    0.00
c3t103d0   0        0         0         0        0.00    0.00
c2t102d0   87       0        44544k      0        0.00    0.00
c3t102d0   0        0         0         0        0.00    0.00
c2t121d0   87       0        44544k      0        0.00    0.00
c3t121d0   0        0         0         0        0.00    0.00
c2t112d0   87       0        44544k      0        0.00    0.00
c3t112d0   0        0         0         0        0.00    0.00
c2t96d0    87       0        44544k      0        0.00    0.00
c3t96d0    0        0         0         0        0.00    0.00
c2t106d0   87       0        44544k      0        0.00    0.00
c3t106d0   0        0         0         0        0.00    0.00
c2t113d0   87       0        44544k      0        0.00    0.00
c3t113d0   0        0         0         0        0.00    0.00

```

```

c2t119d0      87          0      44544k          0          0.00      0.00
c3t119d0       0          0          0          0          0.00      0.00

```

The following command changes the amount of memory that `vxddmpadm` can use to accumulate the statistics:

```
# vxddmpadm iostat start memory=4096
```

The displayed statistics can be filtered by path name, DMP node name, and enclosure name (note that the per-CPU memory has changed following the previous command).

```
# vxddmpadm -u k iostat show pathname=c3t115d0s2
```

```

                                cpu usage = 8132us      per cpu memory = 4096b
                                OPERATIONS              BYTES              AVG TIME (ms)
PATHNAME  READS      WRITES      READS      WRITES      READS      WRITES
c3t115d0s2    0          0          0          0          0.00      0.00

```

```
# vxddmpadm -u k iostat show dmpnodename=c0t0d0s2
```

```

                                cpu usage = 8501us      per cpu memory = 4096b
                                OPERATIONS              BYTES              AVG TIME (ms)
PATHNAME  READS      WRITES      READS      WRITES      READS      WRITES
c0t0d0s2  1088          0      557056k          0          0.00      0.00

```

```
# vxddmpadm -u k iostat show enclosure=Disk
```

```

                                cpu usage = 8626us      per cpu memory = 4096b
                                OPERATIONS              BYTES              AVG TIME (ms)
PATHNAME  READS      WRITES      READS      WRITES      READS      WRITES
c0t0d0s2  1088          0      57056k          0          0.00      0.00

```

You can also specify the number of times to display the statistics and the time interval. Here the incremental statistics for a path are displayed twice with a 2-second interval:

```
# vxddmpadm iostat show dmpnodename=emc_clariion0_342 interval=1 count=2
```

```

                                cpu usage = 164687us     per cpu memory = 409600b
                                OPERATIONS              BLOCKS              AVG TIME (ms)
PATHNAME  READS      WRITES      READS      WRITES      READS      WRITES
c0t5006016041E03B33d6s2  3          0          33          0          0.02      0.00
c0t5006016141E03B33d6s2  3          0           3          0          0.16      0.00
c0t5006016841E03B33d6s2  0          0           0          0          0.00      0.00
c1t5006016041E03B33d6s2  1          0          16          0          0.02      0.00
c1t5006016141E03B33d6s2  2          0           2          0          0.18      0.00
c1t5006016841E03B33d6s2  0          0           0          0          0.00      0.00

```

```

c0t5006016041E03B33d6s2 0      0      0      0      0.00      0.00
c0t5006016141E03B33d6s2 0      0      0      0      0.00      0.00
c0t5006016841E03B33d6s2 0      0      0      0      0.00      0.00
c1t5006016041E03B33d6s2 0      0      0      0      0.00      0.00
c1t5006016141E03B33d6s2 0      0      0      0      0.00      0.00
c1t5006016841E03B33d6s2 0      0      0      0      0.00      0.00

```

Setting the attributes of the paths to an enclosure

You can use the `vxddmpadm setattr` command to set the attributes of the paths to an enclosure or disk array.

```
# vxddmpadm setattr path diskname pathtype=pathtype
```

```
# vxddmpadm setattr path c3t10d0s2 pathtype=nomanual
```

The attributes set for the paths are persistent and are stored in the `/etc/vx/dmppolicy.info` file.

You can set the following attributes:

<code>active</code>	Changes a standby (failover) path to an active path.
<code>nomanual</code>	Restores the original primary or secondary attributes of a path.
<code>nopreferred</code>	Restores the normal priority of a path.
<code>preferred</code> <code>[priority=N]</code>	Specifies a path as preferred, and optionally assigns a priority number to it. If specified, the priority number must be an integer that is greater than or equal to one. Higher priority numbers indicate that a path is able to carry a greater I/O load. Note: Setting a priority for path does not change the I/O policy. The I/O policy must be set independently. See “Specifying the DMP I/O policy” on page 292.
<code>primary</code>	Defines a path as being the primary path for a JBOD disk array.
<code>secondary</code>	Defines a path as being the secondary path for a JBOD disk array.
<code>standby</code>	Marks a standby (failover) path that it is not used for normal I/O scheduling. This path is used if there are no active paths available for I/O.

Displaying the path redundancy level of a device or enclosure under DMP control

Use the `vxddmpadm getddmpnode` command to list the devices with less than the required redundancy level.

To list the devices on a specified enclosure with fewer than a given number of enabled paths, use the following command:

```
# vxddmpadm getddmpnode enclosure=encl_name redundancy=value
```

For example, to list the devices with fewer than 3 enabled paths, use the following command:

```
# vxddmpadm getddmpnode enclosure=EMC_CLARiion0 redundancy=3
```

NAME	STATE	ENCLR-TYPE	PATHS	ENBL	DSBL	ENCLR-NAME
emc_clariion0_162	ENABLED	EMC_CLARiion	3	2	1	emc_clariion0
emc_clariion0_182	ENABLED	EMC_CLARiion	2	2	0	emc_clariion0
emc_clariion0_184	ENABLED	EMC_CLARiion	3	2	1	emc_clariion0
emc_clariion0_186	ENABLED	EMC_CLARiion	2	2	0	emc_clariion0

To display the minimum redundancy level for a particular device, use the `vxddmpadm getattr` command, as follows:

```
# vxddmpadm getattr enclosure|arrayname|arraytype \  
component-name redundancy
```

For example, to show the minimum redundancy level for the enclosure HDS9500-ALUA0:

```
# vxddmpadm getattr enclosure HDS9500-ALUA0 redundancy
```

ENCLR_NAME	DEFAULT	CURRENT
HDS9500-ALUA0	0	4

Specifying the minimum number of active paths for a device or enclosure under DMP control

You can set the minimum redundancy level for a device or an enclosure. The minimum redundancy level is the minimum number of paths that should be active for the device or the enclosure. If the number of paths falls below the minimum redundancy level for the enclosure, a message is sent to the system console and also logged to the DMP log file.

Also, notification is sent to `vxnotify` clients.

The value set for minimum redundancy level is stored in the `dmpolicy.info` file, and is persistent. If no minimum redundancy level is set, the default value is 0.

You can use the `vxddmpadm setattr` command to set the minimum redundancy level.

To specify the minimum number of active paths

- ◆ Use the `vxddmpadm setattr` command with the redundancy attribute as follows:

```
# vxddmpadm setattr enclosure|arrayname|arraytype component-name
  redundancy=value
```

where *value* is the number of active paths.

For example, to set the minimum redundancy level for the enclosure HDS9500-ALUA0:

```
# vxddmpadm setattr enclosure HDS9500-ALUA0 redundancy=2
```

Displaying the DMP I/O policy

To display the current and default settings of the I/O policy for an enclosure, array or array type, use the `vxddmpadm getattr` command.

The following example displays the default and current setting of `iopolicy` for JBOD disks:

```
# vxddmpadm getattr enclosure Disk iopolicy
ENCLR_NAME      DEFAULT      CURRENT
-----
Disk            MinimumQ    Balanced
```

The next example displays the setting of `partitionsize` for the enclosure `enc0`, on which the `balanced` I/O policy with a partition size of 2MB has been set:

```
# vxddmpadm getattr enclosure enc0 partitionsize
```

```
ENCLR_NAME      DEFAULT      CURRENT
-----
enc0            512          4096
```

Specifying the DMP I/O policy

You can use the `vxddmpadm setattr` command to change the I/O policy for distributing I/O load across multiple paths to a disk array or enclosure. You can set policies for an enclosure (for example, `HDS01`), for all enclosures of a particular type (such as `HDS`), or for all enclosures of a particular array type (such as `A/A` for Active/Active, or `A/P` for Active/Passive).

Warning: Starting with release 4.1 of VxVM, I/O policies are recorded in the file `/etc/vx/dmppolicy.info`, and are persistent across reboots of the system.

Do not edit this file yourself.

The following policies may be set:

`adaptive`

This policy attempts to maximize overall I/O throughput from/to the disks by dynamically scheduling I/O on the paths. It is suggested for use where I/O loads can vary over time. For example, I/O from/to a database may exhibit both long transfers (table scans) and short transfers (random look ups). The policy is also useful for a SAN environment where different paths may have different number of hops. No further configuration is possible as this policy is automatically managed by DMP.

In this example, the adaptive I/O policy is set for the enclosure `encl1`:

```
# vxddmpadm setattr enclosure encl1 \
  iopolicy=adaptive
```



```
balanced  
[partitionsizesize]
```

This policy is designed to optimize the use of caching in disk drives and RAID controllers. The size of the cache typically ranges from 120KB to 500KB or more, depending on the characteristics of the particular hardware. During normal operation, the disks (or LUNs) are logically divided into a number of regions (or partitions), and I/O from/to a given region is sent on only one of the active paths. Should that path fail, the workload is automatically redistributed across the remaining paths.

You can use the size argument to the partitionsize attribute to specify the partition size. The partition size in blocks is adjustable in powers of 2 from 2 up to 231. A value that is not a power of 2 is silently rounded down to the nearest acceptable value.

Specifying a partition size of 0 is equivalent to specifying the default partition size.

The default value for the partition size is 512 blocks (256k).

The default value can be changed by adjusting the value of the `dmp_pathswitch_blks_shift` tunable parameter.

See [“DMP tunable parameters”](#) on page 884.

Note: The benefit of this policy is lost if the value is set larger than the cache size.

For example, the suggested partition size for an Hitachi HDS 9960 A/A array is from 32,768 to 131,072 blocks (16MB to 64MB) for an I/O activity pattern that consists mostly of sequential reads or writes.

The next example sets the balanced I/O policy with a partition size of 4096 blocks (2MB) on the enclosure enc0:

```
# vxddmpadm setattr enclosure enc0 \  
iopoly=balanced partitionsize=4096
```

minimumq

This policy sends I/O on paths that have the minimum number of outstanding I/O requests in the queue for a LUN. No further configuration is possible as DMP automatically determines the path with the shortest queue.

The following example sets the I/O policy to `minimumq` for a JBOD:

```
# vxddmpadm setattr enclosure Disk \  
    iopolicy=minimumq
```

This is the default I/O policy for all arrays.

priority

This policy is useful when the paths in a SAN have unequal performance, and you want to enforce load balancing manually. You can assign priorities to each path based on your knowledge of the configuration and performance characteristics of the available paths, and of other aspects of your system.

See [“Setting the attributes of the paths to an enclosure”](#) on page 289.

In this example, the I/O policy is set to `priority` for all SENA arrays:

```
# vxddmpadm setattr arrayname SENA \  
    iopolicy=priority
```

round-robin

This policy shares I/O equally between the paths in a round-robin sequence. For example, if there are three paths, the first I/O request would use one path, the second would use a different path, the third would be sent down the remaining path, the fourth would go down the first path, and so on. No further configuration is possible as this policy is automatically managed by DMP.

The next example sets the I/O policy to `round-robin` for all Active/Active arrays:

```
# vxddmpadm setattr arraytype A/A \  
    iopolicy=round-robin
```

singleactive

This policy routes I/O down the single active path. This policy can be configured for A/P arrays with one active path per controller, where the other paths are used in case of failover. If configured for A/A arrays, there is no load balancing across the paths, and the alternate paths are only used to provide high availability (HA). If the current active path fails, I/O is switched to an alternate active path. No further configuration is possible as the single active path is selected by DMP.

The following example sets the I/O policy to `singleactive` for JBOD disks:

```
# vxddmpadm setattr arrayname Disk \  
iopolicy=singleactive
```

Scheduling I/O on the paths of an Asymmetric Active/Active array

You can specify the `use_all_paths` attribute in conjunction with the `adaptive`, `balanced`, `minimumq`, `priority` and `round-robin` I/O policies to specify whether I/O requests are to be scheduled on the secondary paths in addition to the primary paths of an Asymmetric Active/Active (A/A-A) array. Depending on the characteristics of the array, the consequent improved load balancing can increase the total I/O throughput. However, this feature should only be enabled if recommended by the array vendor. It has no effect for array types other than A/A-A.

The following example command sets the `balanced` I/O policy on the enclosure `enc0`, and allows scheduling of I/O requests on the secondary paths.

In this example, the partition size is 4096 blocks (2MB).

```
# vxddmpadm setattr enclosure enc0 iopolicy=balanced \  
partitionsize=4096 use_all_paths=yes
```

The default setting for this attribute is `use_all_paths=no`.

You can display the current setting for `use_all_paths` for an enclosure, arrayname or arraytype. To do this, specify the `use_all_paths` option to the `vxddmpadm gettattr` command.

```
# vxddmpadm gettattr enclosure HDS9500-ALUA0 use_all_paths
```

ENCLR_NAME	DEFAULT	CURRENT
=====		
HDS9500-ALUA0	no	yes

The `use_all_paths` attribute only applies to A/A-A arrays. For other arrays, the above command displays the message:

```
Attribute is not applicable for this array.
```

Example of applying load balancing in a SAN

This example describes how to configure load balancing in a SAN environment where there are multiple primary paths to an Active/Passive device through several SAN switches. As shown in the sample output from the `vxdisk list` command, the device has eight primary paths:

```
# vxdisk list c3t2d15s2

Device: c3t2d15s2
.
.
.

numpaths: 8
c2t0d15s2 state=enabled type=primary
c2t1d15s2 state=enabled type=primary
c3t1d15s2 state=enabled type=primary
c3t2d15s2 state=enabled type=primary
c4t2d15s2 state=enabled type=primary
c4t3d15s2 state=enabled type=primary
c5t3d15s2 state=enabled type=primary
c5t4d15s2 state=enabled type=primary
```

In addition, the device is in the enclosure `ENC0`, belongs to the disk group `mydg`, and contains a simple concatenated volume `myvol1`.

The first step is to enable the gathering of DMP statistics:

```
# vxdkmpadm iostat start
```

Next, use the `dd` command to apply an input workload from the volume:

```
# dd if=/dev/vx/rdisk/mydg/myvol1 of=/dev/null &
```

Use the `vxdkmpadm iostat` command to display the DMP statistics for the device. The sample output shows that all I/O is being directed to one path.

```
# vxdkmpadm iostat show dmpnodename=c3t2d15s2 interval=5 count=2
.
.
.
```

cpu usage = 11294us per cpu memory = 32768b

PATHNAME	OPERATIONS		KBYTES		AVG TIME (ms)	
	READS	WRITES	READS	WRITES	READS	WRITES
c2t0d15s2	0	0	0	0	0.00	0.00
c2t1d15s2	0	0	0	0	0.00	0.00
c3t1d15s2	0	0	0	0	0.00	0.00
c3t2d15s2	0	0	0	0	0.00	0.00
c4t2d15s2	0	0	0	0	0.00	0.00
c4t3d15s2	0	0	0	0	0.00	0.00
c5t3d15s2	0	0	0	0	0.00	0.00
c5t4d15s2	10986	0	5493	0	0.41	0.00

Use the `vxddmpadm getattr enclosure ENC0 iopolicy` command to display the I/O policy for the enclosure that contains the device:

```
# vxddmpadm getattr enclosure ENC0 iopolicy
```

ENCLR_NAME	DEFAULT	CURRENT
=====		
ENC0	MinimumQ	Single-Active

The above output shows that the policy for the enclosure is set to `singleactive`, which explains why all the I/O is taking place on one path.

To balance the I/O load across the multiple primary paths, set the policy to `round-robin` as shown here:

```
# vxddmpadm setattr enclosure ENC0 iopolicy=round-robin
# vxddmpadm getattr enclosure ENC0 iopolicy
```

ENCLR_NAME	DEFAULT	CURRENT
=====		
ENC0	MinimumQ	Round-Robin

Reset the DMP statistics:

```
# vxddmpadm iostat reset
```

With the workload still running, you can see the effect of changing the I/O policy to balance the load across the primary paths.

```
# vxddmpadm iostat show dmpnodename=c3t2d15s2 interval=5 count=2
.
.
.
```

cpu usage = 14403us per cpu memory = 32768b

	OPERATIONS	KBYTES	AVG TIME (ms)

PATHNAME	READS	WRITES	READS	WRITES	READS	WRITES
c2t0d15s2	2041	0	1021	0	0.39	0.00
c2t1d15s2	1894	0	947	0	0.39	0.00
c3t1d15s2	2008	0	1004	0	0.39	0.00
c3t2d15s2	2054	0	1027	0	0.40	0.00
c4t2d15s2	2171	0	1086	0	0.39	0.00
c4t3d15s2	2095	0	1048	0	0.39	0.00
c5t3d15s2	2073	0	1036	0	0.39	0.00
c5t4d15s2	2042	0	1021	0	0.39	0.00

Enter the following command to return the enclosure to the single active I/O policy.

```
# vxdkmpadm setattr enclosure ENC0 iopolicy=singleactive
```

Disabling I/O for paths, controllers or array ports

Disabling I/O through a path, HBA controller or array port prevents DMP from issuing I/O requests through the specified path, or the paths that are connected to the specified controller or array port. The command blocks until all pending I/O requests issued through the paths are completed.

Note: From release 5.0 of VxVM, this operation is supported for controllers that are used to access disk arrays on which cluster-shareable disk groups are configured.

Before detaching a system board, stop all I/O to the HBA controllers that are located on the board. To do this, execute the `vxdkmpadm disable` command, and then run the Dynamic Reconfiguration (DR) facility provided by Oracle.

DMP does not support the operation to disable I/O for the controllers that use Third-Party Drivers (TPD) for multi-pathing.

To disable I/O for one or more paths, use the following command:

```
# vxdkmpadm [-c|-f] disable path=path_name1[,path_name2,path_nameN]
```

To disable I/O for the paths connected to one or more HBA controllers, use the following command:

```
# vxdkmpadm [-c|-f] disable ctrl=ctrl_name1[,ctrl_name2,ctrl_nameN]
```

To disable I/O for the paths connected to an array port, use one of the following commands:

```
# vxdkmpadm [-c|-f] disable enclosure=enclr_name portid=array_port_ID
# vxdkmpadm [-c|-f] disable pwwn=array_port_WWN
```

where the array port is specified either by the enclosure name and the array port ID, or by the array port's worldwide name (WWN) identifier.

The following examples show how to disable I/O on an array port:

```
# vxddpdm disable enclosure=HDS9500V0 portid=1A
# vxddpdm disable pwwn=20:00:00:E0:8B:06:5F:19
```

To disable I/O for a particular path, specify both the controller and the portID, which represent the two ends of the fabric:

```
# vxddpdm [-c|-f] disable ctrl=ctrl_name enclosure=enclr_name \
portid=array_port_ID
```

You can use the `-c` option to check if there is only a single active path to the disk. If so, the `disable` command fails with an error message unless you use the `-f` option to forcibly disable the path.

The `disable` operation fails if it is issued to a controller that is connected to the root disk through a single path, and there are no root disk mirrors configured on alternate paths. If such mirrors exist, the command succeeds.

Enabling I/O for paths, controllers or array ports

Enabling a controller allows a previously disabled path, HBA controller or array port to accept I/O again. This operation succeeds only if the path, controller or array port is accessible to the host, and I/O can be performed on it. When connecting Active/Passive disk arrays, the `enable` operation results in failback of I/O to the primary path. The `enable` operation can also be used to allow I/O to the controllers on a system board that was previously detached.

Note: This operation is supported for controllers that are used to access disk arrays on which cluster-shareable disk groups are configured.

DMP does not support the operation to enable I/O for the controllers that use Third-Party Drivers (TPD) for multi-pathing.

To enable I/O for one or more paths, use the following command:

```
# vxddpdm enable path=path_name1[,path_name2,path_nameN]
```

To enable I/O for the paths connected to one or more HBA controllers, use the following command:

```
# vxddpdm enable ctrl=ctrl_name1[,ctrl_name2,ctrl_nameN]
```

To enable I/O for the paths connected to an array port, use one of the following commands:

```
# vxddmpadm enable enclosure=enclr_name portid=array_port_ID
# vxddmpadm enable pwwn=array_port_WWN
```

where the array port is specified either by the enclosure name and the array port ID, or by the array port’s worldwide name (WWN) identifier.

The following are examples of using the command to enable I/O on an array port:

```
# vxddmpadm enable enclosure=HDS9500V0 portid=1A
# vxddmpadm enable pwwn=20:00:00:E0:8B:06:5F:19
```

To enable I/O for a particular path, specify both the controller and the portID, which represent the two ends of the fabric:

```
# vxddmpadm enable ctrl=ctrl_name enclosure=enclr_name \
  portid=array_port_ID
```

Renaming an enclosure

The `vxddmpadm setattr` command can be used to assign a meaningful name to an existing enclosure, for example:

```
# vxddmpadm setattr enclosure enc0 name=GRP1
```

This example changes the name of an enclosure from `enc0` to `GRP1`.

Note: The maximum length of the enclosure name prefix is 23 characters.

The following command shows the changed name:

```
# vxddmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS
=====			
other0	OTHER	OTHER_DISKS	CONNECTED
jbod0	X1	X1_DISKS	CONNECTED
GRP1	ACME	60020f20000001a90000	CONNECTED

Configuring DMP's response to I/O failures

You can configure how DMP responds to failed I/O requests on the paths to a specified enclosure, disk array name, or type of array. By default, DMP is configured to retry a failed I/O request up to five times for a single path.

To display the current settings for handling I/O request failures that are applied to the paths to an enclosure, array name or array type, use the `vxddmpadm getattr` command.

See [“Displaying recovery option values”](#) on page 304.

To set a limit for the number of times that DMP attempts to retry sending an I/O request on a path, use the following command:

```
# vxddmpadm setattr \  
  {enclosure enc-name|arrayname name|arraytype type} \  
  recoveryoption=fixedretry retrycount=n
```

The value of the argument to `retrycount` specifies the number of retries to be attempted before DMP reschedules the I/O request on another available path, or fails the request altogether.

As an alternative to specifying a fixed number of retries, you can specify the amount of time DMP allows for handling an I/O request. If the I/O request does not succeed within that time, DMP fails the I/O request. To specify an `iotimeout` value, use the following command:

```
# vxddmpadm setattr \  
  {enclosure enc-name|arrayname name|arraytype type} \  
  recoveryoption=timebound iotimeout=seconds
```

The default value of `iotimeout` is 300 seconds. For some applications such as Oracle, it may be desirable to set `iotimeout` to a larger value. The `iotimeout` value for DMP should be greater than the I/O service time of the underlying operating system layers.

Note: The `fixedretry` and `timebound` settings are mutually exclusive.

The following example configures time-bound recovery for the enclosure `enc0`, and sets the value of `iotimeout` to 360 seconds:

```
# vxddmpadm setattr enclosure enc0 recoveryoption=timebound \  
  iotimeout=360
```

The next example sets a fixed-retry limit of 10 for the paths to all Active/Active arrays:

```
# vxddmpadm setattr arraytype A/A recoveryoption=fixedretry \  
    retrycount=10
```

Specifying `recoveryoption=default` resets DMP to the default settings corresponding to `recoveryoption=fixedretry retrycount=5`, for example:

```
# vxddmpadm setattr arraytype A/A recoveryoption=default
```

The above command also has the effect of configuring I/O throttling with the default settings.

See [“Configuring the DMP I/O throttling mechanism”](#) on page 302.

Note: The response to I/O failure settings is persistent across reboots of the system.

Configuring the DMP I/O throttling mechanism

By default, DMP is configured with I/O throttling turned off for all paths. To display the current settings for I/O throttling that are applied to the paths to an enclosure, array name or array type, use the `vxddmpadm getattr` command.

See [“Displaying recovery option values”](#) on page 304.

If enabled, I/O throttling imposes a small overhead on CPU and memory usage because of the activity of the statistics-gathering daemon. If I/O throttling is disabled, the daemon no longer collects statistics, and remains inactive until I/O throttling is re-enabled.

To turn off I/O throttling, use the following form of the `vxddmpadm setattr` command:

```
# vxddmpadm setattr \  
    {enclosure enc-name|arrayname name|arraytype type} \  
    recoveryoption=nothrottle
```

The following example shows how to disable I/O throttling for the paths to the enclosure `enc0`:

```
# vxddmpadm setattr enclosure enc0 recoveryoption=nothrottle
```

The `vxddmpadm setattr` command can be used to enable I/O throttling on the paths to a specified enclosure, disk array name, or type of array:

```
# vxddmpadm setattr \  
  {enclosure enc-name|arrayname name|arraytype type}\  
  recoveryoption=throttle [iotimeout=seconds]
```

If the `iotimeout` attribute is specified, its argument specifies the time in seconds that DMP waits for an outstanding I/O request to succeed before invoking I/O throttling on the path. The default value of `iotimeout` is 10 seconds. Setting `iotimeout` to a larger value potentially causes more I/O requests to become queued up in the SCSI driver before I/O throttling is invoked.

The following example sets the value of `iotimeout` to 60 seconds for the enclosure `enc0`:

```
# vxddmpadm setattr enclosure enc0 recoveryoption=throttle \  
  iotimeout=60
```

Specify `recoveryoption=default` to reset I/O throttling to the default settings, as follows:

```
# vxddmpadm setattr arraytype A/A recoveryoption=default
```

The above command configures the default behavior, corresponding to `recoveryoption=nothrottle`. The above command also configures the default behavior for the response to I/O failures.

See [“Configuring DMP's response to I/O failures”](#) on page 301.

Note: The I/O throttling settings are persistent across reboots of the system.

Configuring DMP's Low Impact Path Probing

See [“Low Impact Path Probing \(LIPP\)”](#) on page 178.

The Low Impact Path Probing (LIPP) feature can be turned on or off using the `vxddmpadm settune` command:

```
# vxddmpadm settune dmp_low_impact_probe=[on|off]
```

Path probing will be optimized by probing a subset of paths connected to the same HBA and array port. The size of the subset of paths can be controlled by the `dmp_probe_threshold` tunable. The default value is set to 5.

```
# vxddmpadm settune dmp_probe_threshold=N
```

Configuring DMP's Subpaths Failover Groups (SFG)

See “Subpaths Failover Group (SFG)” on page 178.

The Subpaths Failover Groups (SFG) feature can be turned on or off using the tunable `dmp_sfg_threshold`.

To turn off the feature, set the tunable `dmp_sfg_threshold` value to 0:

```
# vxddmpadm settune dmp_sfg_threshold=0
```

To turn on the feature, set the `dmp_sfg_threshold` value to the required number of path failures which triggers SFG. The default is 1.

```
# vxddmpadm settune dmp_sfg_threshold=N
```

The default value of the tunable is “1” which represents that the feature is on.

To see the Subpaths Failover Groups ID, use the following command:

```
# vxddmpadm getportids {ctlr=ctlr_name | dmpnodename=dmp_device_name \
| enclosure=enclr_name | path=path_name}
```

Displaying recovery option values

To display the current settings for handling I/O request failures that are applied to the paths to an enclosure, array name or array type, use the following command:

```
# vxddmpadm getattr \
{enclosure enc-name|arrayname name|arraytype type} \
recoveryoption
```

The following example shows the `vxddmpadm getattr` command being used to display the `recoveryoption` option values that are set on an enclosure.

```
# vxddmpadm getattr enclosure HDS9500-ALUA0 recoveryoption
ENCLR-NAME      RECOVERY-OPTION  DEFAULT[VAL]    CURRENT[VAL]
=====
HDS9500-ALUA0   Throttle         Nothrottle[0]   Timebound[60]
HDS9500-ALUA0   Error-Retry      Fixed-Retry[5]  Timebound[20]
```

This shows the default and current policy options and their values.

[Table 11-3](#) summarizes the possible recovery option settings for retrying I/O after an error.

Table 11-3 Recovery options for retrying I/O after an error

Recovery option	Possible settings	Description
recoveryoption=fixedretry	Fixed-Retry (retrycount)	DMP retries a failed I/O request for the specified number of times if I/O fails.
recoveryoption=timebound	Timebound (iotimeout)	DMP retries a failed I/O request for the specified time in seconds if I/O fails.

[Table 11-4](#) summarizes the possible recovery option settings for throttling I/O.

Table 11-4 Recovery options for I/O throttling

Recovery option	Possible settings	Description
recoveryoption=nothrottle	None	I/O throttling is not used.
recoveryoption=throttle	Timebound (iotimeout)	DMP throttles the path if an I/O request does not return within the specified time in seconds.

Configuring DMP path restoration policies

DMP maintains a kernel thread that re-examines the condition of paths at a specified interval. The type of analysis that is performed on the paths depends on the checking policy that is configured.

Note: The DMP path restoration thread does not change the disabled state of the path through a controller that you have disabled using `vxddmpadm disable`.

When configuring DMP path restoration policies, you must stop the path restoration thread, and then restart it with new attributes.

See [“Stopping the DMP path restoration thread”](#) on page 307.

Use the `vxddmpadm settune dmp_restore_policy` command to configure one of the following restore policies. The policy will remain in effect until the restore thread is stopped or the values are changed using `vxddmpadm settune` command.

- `check_all`

The path restoration thread analyzes all paths in the system and revives the paths that are back online, as well as disabling the paths that are inaccessible. The command to configure this policy is:

```
# vxddmpadm settune dmp_restore_policy=check_all
```

- `check_alterate`

The path restoration thread checks that at least one alternate path is healthy. It generates a notification if this condition is not met. This policy avoids inquiry commands on all healthy paths, and is less costly than `check_all` in cases where a large number of paths are available. This policy is the same as `check_all` if there are only two paths per DMP node. The command to configure this policy is:

```
# vxddmpadm settune dmp_restore_policy=check_alterate
```

- `check_disabled`

This is the default path restoration policy. The path restoration thread checks the condition of paths that were previously disabled due to hardware failures, and revives them if they are back online. The command to configure this policy is:

```
# vxddmpadm settune dmp_restore_policy=check_disabled
```

- `check_periodic`

The path restoration thread performs `check_all` once in a given number of cycles, and `check_disabled` in the remainder of the cycles. This policy may lead to periodic slowing down (due to `check_all`) if a large number of paths are available. The command to configure this policy is:

```
# vxddmpadm settune dmp_restore_policy=check_periodic
```

The default number of cycles between running the `check_all` policy is 10.

The `dmp_restore_interval` tunable parameter specifies how often the path restoration thread examines the paths. For example, the following command sets the polling interval to 400 seconds:

```
# vxddmpadm settune dmp_restore_interval=400
```

The settings are immediately applied and are persistent across reboots. Use the `vxddmpadm gettune` to view the current settings.

See “DMP tunable parameters” on page 884.

If the `vxddmpadm start restore` command is given without specifying a policy or interval, the path restoration thread is started with the persistent policy and interval settings previously set by the administrator with the `vxddmpadm settune` command. If the administrator has not set a policy or interval, the system defaults are used.

The system default restore policy is `check_disabled`. The system default interval is 300 seconds.

Warning: Decreasing the interval below the system default can adversely affect system performance.

Stopping the DMP path restoration thread

Use the following command to stop the DMP path restoration thread:

```
# vxddmpadm stop restore
```

Warning: Automatic path failback stops if the path restoration thread is stopped.

Displaying the status of the DMP path restoration thread

Use the `vxddmpadm gettune` command to display the tunable parameter values that show the status of the DMP path restoration thread. These tunables include:

`dmp_restore_state` the status of the automatic path restoration kernel thread.

`dmp_restore_interval` the polling interval for the DMP path restoration thread.

`dmp_restore_policy` the policy that DMP uses to check the condition of paths.

To display the status of the DMP path restoration thread

◆ Use the following commands:

```
# vxddmpadm gettune dmp_restore_state
```

```
# vxddmpadm gettune dmp_restore_interval
```

```
# vxddmpadm gettune dmp_restore_policy
```

Configuring array policy modules

An array policy module (APM) is a dynamically loadable kernel module (plug-in for DMP) for use in conjunction with an array. An APM defines array-specific procedures and commands to:

- Select an I/O path when multiple paths to a disk within the array are available.
- Select the path failover mechanism.

- Select the alternate path in the case of a path failure.
- Put a path change into effect.
- Respond to SCSI reservation or release requests.

DMP supplies default procedures for these functions when an array is registered. An APM may modify some or all of the existing procedures that are provided by DMP or by another version of the APM.

You can use the following command to display all the APMs that are configured for a system:

```
# vxddmpadm listapm all
```

The output from this command includes the file name of each module, the supported array type, the APM name, the APM version, and whether the module is currently loaded and in use. To see detailed information for an individual module, specify the module name as the argument to the command:

```
# vxddmpadm listapm module_name
```

To add and configure an APM, use the following command:

```
# vxddmpadm -a cfgapm module_name [attr1=value1 \  
    [attr2=value2 ...]]
```

The optional configuration attributes and their values are specific to the APM for an array. Consult the documentation that is provided by the array vendor for details.

Note: By default, DMP uses the most recent APM that is available. Specify the `-u` option instead of the `-a` option if you want to force DMP to use an earlier version of the APM. The current version of an APM is replaced only if it is not in use.

Specifying the `-r` option allows you to remove an APM that is not currently loaded:

```
# vxddmpadm -r cfgapm module_name
```

See the `vxddmpadm(1M)` manual page.

DMP coexistence with native multi-pathing

Dynamic Multi-Pathing (DMP) supports using multi-pathing with raw devices. The `dmp_native_multipathing` tunable controls the behavior. If the `dmp_native_multipathing` tunable is set to on, DMP intercepts I/O requests, operations such as open, close, and ioctl sent on the raw device path.

If the `dmp_native_multipathing` tunable is set to off, these requests are sent directly to the raw device. In A/PF arrays, the format command on Solaris platform does not show the extra attributes (like vendor ID, product ID and geometry information) of the passive paths. To avoid this issue, enable the `dmp_native_multipathing` tunable. DMP intercepts the request and routes it on the primary path.

For A/P arrays, turning on the `dmp_native_multipathing` feature enables the commands to succeed without trespassing. The feature has no benefit for A/A or A/A-A arrays.

DMP native multi-pathing should not be enabled if one of the following tools are already managing multi-pathing:

- EMC PowerPath
- Sun StorEdge Traffic Manager (also called MPxIO)

If EMC PowerPath is installed first, the command to set `dmp_native_multipathing` to on fails. If VxVM is installed first, ensure that `dmp_native_multipathing` is set to off before installing EMC PowerPath.

Dynamic reconfiguration of devices

This chapter includes the following topics:

- [About online dynamic reconfiguration](#)
- [Reconfiguring a LUN online that is under DMP control](#)
- [Removing LUNs dynamically from an existing target ID](#)
- [Adding new LUNs dynamically to a new target ID](#)
- [About detecting target ID reuse if the operating system device tree is not cleaned up](#)
- [Scanning an operating system device tree after adding or removing LUNs](#)
- [Cleaning up the operating system device tree after removing LUNs](#)
- [Dynamic LUN expansion](#)
- [Replacing a host bus adapter on an M5000 server managed by DMP](#)
- [Upgrading the array controller firmware online](#)

About online dynamic reconfiguration

You can perform the following kinds of online dynamic reconfigurations:

- Reconfiguring a LUN online that is under DMP control
See [“Reconfiguring a LUN online that is under DMP control”](#) on page 311.
- Updating the array controller firmware, also known as a nondisruptive upgrade
See [“Upgrading the array controller firmware online”](#) on page 323.

- Replacing a host bus adapter (HBA) online
See [“Replacing a host bus adapter on an M5000 server managed by DMP”](#) on page 319.

Reconfiguring a LUN online that is under DMP control

System administrators and storage administrators may need to modify the set of LUNs provisioned to a server. You can change the LUN configuration dynamically, without performing a reconfiguration reboot on the host.

Dynamic LUN reconfigurations require array configuration commands, operating system commands, and Veritas Volume manager commands. To complete the operations correctly, you must issue the commands in the proper sequence on the host.

The operations are as follows:

- Dynamic LUN removal from an existing target ID
See [“Removing LUNs dynamically from an existing target ID”](#) on page 311.
- Dynamic new LUN addition to a new target ID
See [“Adding new LUNs dynamically to a new target ID”](#) on page 314.
- Dynamic LUN expansion
See [“Dynamic LUN expansion”](#) on page 318.

Removing LUNs dynamically from an existing target ID

In this case, a group of LUNs is unmapped from the host HBA ports and an operating system device scan is issued. To add subsequent LUNs seamlessly, perform additional steps to clean up the operating system device tree.

The high-level procedure and the SFCFSA commands are generic. However, the operating system commands may vary depending on the Solaris version. For example, the following procedure uses Solaris 10 with the Leadville stack.

To remove LUNs dynamically from an existing target ID

- 1 Prior to any dynamic reconfiguration, ensure that the `dmp_cache_open` tunable is set to `on`. This setting is the default.

```
# vxddm padm gettune dmp_cache_open
```

If the tunable is set to `off`, set the `dmp_cache_open` tunable to `on`.

```
# vxddm padm settune dmp_cache_open=on
```

- 2 Identify which LUNs to remove from the host. Do one of the following:
 - Use Storage Array Management to identify the Array Volume ID (AVID) for the LUNs.
 - If the array does not report the AVID, use the LUN index.
- 3 For LUNs under VxVM, perform the following steps:

- Evacuate the data from the LUNs using the `vxevac` command.

See the `vxevac(1M)` online manual page.

After the data has been evacuated, enter the following command to remove the LUNs from the disk group:

```
# vxddg -g diskgroup rmdisk da-name
```

- If the data has not been evacuated and the LUN is part of a subdisk or disk group, enter the following command to remove the LUNs from the disk group. If the disk is part of a shared disk group, you must use the `-k` option to force the removal.

```
# vxddg -g diskgroup -k rmdisk da-name
```

- 4 For LUNs that are in use by ZFS, export or destroy the zpool.
- 5 Using the AVID or LUN index, use Storage Array Management to unmap or unmask the LUNs you identified in step 2.
- 6 Remove the LUNs from the `vxddisk` list. Enter the following command on all nodes in a cluster:

```
# vxddisk rm da-name
```

This is a required step. If you do not perform this step, the DMP device tree shows ghost paths.

- 7 Clean up the Solaris SCSI device tree for the devices that you removed in step 6.

See [“Cleaning up the operating system device tree after removing LUNs”](#) on page 316.

This step is required. You must clean up the operating system SCSI device tree to release the SCSI target ID for reuse if a new LUN is added to the host later.

- 8 Scan the operating system device tree.

See [“Scanning an operating system device tree after adding or removing LUNs”](#) on page 316.

- 9 Use SFCFSHA to perform a device scan. You must perform this operation on all nodes in a cluster. Enter one of the following commands:

- `# vxdctl enable`
- `# vxdisk scandisks`

- 10 Refresh the DMP device name database using the following command:

- `# vxddladm assign names`

- 11 Verify that the LUNs were removed cleanly by answering the following questions:

- Is the device tree clean?
Verify that the operating system metanodes are removed from the `/dev` directory.
- Were all the appropriate LUNs removed?
Use the DMP disk reporting tools such as the `vxdisk list` command output to determine if the LUNs have been cleaned up successfully.
- Is the `vxdisk list` output correct?
Verify that the `vxdisk list` output shows the correct number of paths and does not include any ghost disks.

If the answer to any of these questions is "No," return to step 5 and perform the required steps.

If the answer to all of the questions is "Yes," the LUN remove operation is successful.

Adding new LUNs dynamically to a new target ID

In this case, a new group of LUNs is mapped to the host via multiple HBA ports. An operating system device scan is issued for the LUNs to be recognized and added to DMP control.

The high-level procedure and the SFCFSHA commands are generic. However, the operating system commands may vary depending on the Solaris version. For example, the following procedure uses Solaris 10 with the Leadville stack.

To add new LUNs dynamically to a new target ID

- 1 Prior to any dynamic reconfiguration, ensure that the `dmp_cache_open` tunable is set to `on`. This setting is the default.

```
# vxddmpadm gettune dmp_cache_open
```

If the tunable is set to `off`, set the `dmp_cache_open` tunable to `on`.

```
# vxddmpadm settune dmp_cache_open=on
```

- 2 If DMP co-exists with EMC PowerPath, make sure the `dmp_monitor_osevent` parameter is set to `off`. The `vxesd` daemon will not monitor operating system events.

If you install DMP on a system that already has PowerPath installed, DMP sets the `dmp_monitor_osevent` to `off` by default.

```
# vxddmpadm gettune dmp_monitor_osevent
```

If required, turn off the `dmp_monitor_osevent` parameter explicitly:

```
# vxddmpadm settune dmp_monitor_osevent=off
```

- 3 Identify which LUNs to add to the host. Do one of the following:
 - Use Storage Array Management to identify the Array Volume ID (AVID) for the LUNs.
 - If the array does not report the AVID, use the LUN index.
- 4 Map/mask the LUNs to the new target IDs on multiple hosts.
- 5 Scan the operating system device.

See [“Scanning an operating system device tree after adding or removing LUNs”](#) on page 316.

Repeat step 3 and step 4 until you see that all the LUNs have been added.

About detecting target ID reuse if the operating system device tree is not cleaned up

- 6 Use SFCFSHA to perform a device scan. You must perform this operation on all nodes in a cluster. Enter one of the following commands:

```
■ # vxdctl enable

■ # vxdisk scandisks
```

- 7 Refresh the DMP device name database using the following command:

```
# vxddladm assign names
```

- 8 Verify that the LUNs were added correctly by answering the following questions:

- Do the newly provisioned LUNs appear in the `vxdisk list` output?
- Are the configured paths present for each LUN?

If the answer to any of these questions is "No," return to step 3 and begin the procedure again.

If the answer to all of the questions is "Yes," the LUNs have been successfully added. You can now add the LUNs to a disk group, create new volumes, or grow existing volumes.

If the `dmp_native_support` tunable is set to ON and the new LUN does not have a VxVM label or is not claimed by a TPD driver then the LUN is available for use by LVM.

About detecting target ID reuse if the operating system device tree is not cleaned up

If you try to reprovision a LUN or set of LUNs whose previously-valid operating system device entries are not cleaned up, the following messages are displayed. Also, DMP reconfiguration during the DMP device scan and DMP reconfiguration are temporarily inhibited.

See [“Cleaning up the operating system device tree after removing LUNs”](#) on page 316.

```
VxVM vxdisk ERROR V-5-1-14519 Data Corruption Protection Activated
- User Corrective Action Needed
```

```
VxVM vxdisk INFO V-5-1-14521 To recover, first ensure that the OS
device tree is up to date (requires OS specific commands).
```

```
VxVM vxdisk INFO V-5-1-14520 Then, execute 'vxdisk rm' on the
following devices before reinitiating device discovery. <DA names>
```

The message above indicates that a new LUN is trying to reuse the target ID of an older LUN. The device entries have not been cleaned, so the new LUN cannot use the target ID. Until the operating system device tree is cleaned up, DMP prevents this operation.

Scanning an operating system device tree after adding or removing LUNs

After you add or remove LUNs, scan the operating system device tree to verify that the operation completed successfully.

The operating system commands may vary, depending on the Solaris version. The following procedure uses Solaris 10 with the Leadville stack.

To scan an operating system device tree after adding or removing LUNs

- 1 Enter the following command:

```
# cfgadm -c configure c2
```

where `c2` is the controller ID 2.

- 2 Enter the following command:

```
# devfsadm -Cv
```

Cleaning up the operating system device tree after removing LUNs

After you remove LUNs, you must clean up the operating system device tree.

The operating system commands may vary, depending on the Solaris version. The following procedure uses Solaris 10 with the Leadville stack. If any of these steps do not produce the desired result, contact Oracle support.

To clean up the operating system device tree after removing LUNs

- 1 Run the `format` command. In the command output, a device that has been removed includes the text `<drive not available>`.

```
413. c3t5006048ACAFE4A7Cd252 <drive not available>
/pci@1d,700000/SUNW,qlc@1,1/fp@0,0/ssd@w5006048acafe4a7c,fc
```

- 2 Use Storage Array Management or the command line to unmap the LUNs. After they are unmapped, Solaris indicates the devices are either `unusable` or `failing`.

```
# cfgadm -al -o show_SCSI_LUN | grep -i unusable
c2::5006048acafe4a73,256 disk connected configured unusable
c3::5006048acafe4a7c,255 disk connected configured unusable
# cfgadm -al -o show_SCSI_LUN | grep -i failing
c2::5006048acafe4a73,71 disk connected configured failing
c3::5006048acafe4a7c,252 disk connected configured failing
```

See [“Reconfiguring a LUN online that is under DMP control”](#) on page 311.

- 3 If the output indicates the LUNs are `failing`, you must force an LIP on the HBA.

```
# luxadm -e forcelp /devices/pci@1d,700000/SUNW,qlc@1,1/fp
@0,0:devctl
```

This operation probes the targets again, so that output indicates the devices are `unusable`. To remove a device from the operating system device tree, it must be `unusable`.

- 4 Remove the device from the `cfgadm` database. On the HBA, enter the following commands:

```
# cfgadm -c unconfigure -o unusable_SCSI_LUN c2::5006048acafe4a73
# cfgadm -c unconfigure -o unusable_SCSI_LUN c3::5006048acafe4a7c
```

- 5 To verify that the LUNs have been removed, repeat step 2.
- 6 Clean up the device tree. The following command removes the `/dev/dsk` links to `/devices`.

```
# devfsadm -Cv
```

Dynamic LUN expansion

Many modern disk arrays allow existing LUNs to be resized. The Veritas Volume Manager (VxVM) supports dynamic LUN expansion, by providing a facility to update disk headers and other VxVM structures to match a new LUN size. The device must have a SCSI interface that is presented by a smart switch, smart array or RAID controller.

Resizing should only be performed on LUNs 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).

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

For disks with the VxVM `cdsdisk` layout, disks larger than 1 TB in size have a different internal layout than disks smaller than 1 TB. Therefore, resizing a `cdsdisk` disk from less than 1 TB to greater than 1 TB requires special care. If the disk has the VxVM disk group has only one disk, which has the `cdsdisk` layout, you must add a second disk (of any size) to the disk group prior to performing the `vxdisk resize` command on the original disk. You can remove the second disk from the disk group after the resize operation has completed.

Use the following form of the `vxdisk` command to make VxVM aware of the new size of a LUN that has been resized:

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

If a disk media name rather than a disk access name is specified, a disk group name is required. Specify a disk group with the `-g` option or configure a default disk group. If the default disk group is not configured, the above command generates an error message.

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.

Any volumes on the device should only be grown after the LUN itself has first been grown.

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

Before shrinking a LUN, first shrink any volumes on the LUN or move those volumes off of the LUN. Then, resize the device using `vxdisk resize`. Finally, resize the LUN itself using the storage array's management utilities. 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. Note the following exception. For disks with the VxVM `cdsdisk` layout, disks larger than 1 TB in size have a different internal layout than disks smaller than 1 TB. Therefore, resizing a `cdsdisk` disk from less than 1 TB to greater than 1 TB requires special care if the disk group only has one disk. In this case, you must add a second disk (of any size) to the disk group prior to performing the `vxdisk resize` command on the original disk. You can remove the second disk from the disk group after the resize operation has completed.

Caution: 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.

Replacing a host bus adapter on an M5000 server managed by DMP

This section contains the procedure to replace an online host bus adapter (HBA) when DMP is managing multi-pathing in a Cluster File System (CFS) cluster. The HBA World Wide Port Name (WWPN) changes when the HBA is replaced. Following are the prerequisites to replace an online host bus adapter:

- A single node or two or more node CFS or RAC cluster.
- I/O running on CFS file system.
- An M5000 server with at least two HBAs in separate PCIe slots and recommended Solaris patch level for HBA replacement.

To replace an online host bus adapter on an M5000 server

- 1 Identify the HBAs on the M5000 server. For example, to identify Emulex HBAs, enter the following command:

```
/usr/platform/sun4u/sbin/prtdiag -v | grep emlx
00 PCIE 0 2, fc20, 10df 119, 0, 0 okay 4,
4 SUNW,emlxs-pci10df,fc20 LPe 11002-S
/pci@0,600000/pci@0/pci@9/SUNW,emlxs@0
00 PCIE 0 2, fc20, 10df 119, 0, 1 okay 4,
4 SUNW,emlxs-pci10df,fc20 LPe 11002-S
/pci@0,600000/pci@0/pci@9/SUNW,emlxs@0,1
00 PCIE 3 2, fc20, 10df 2, 0, 0 okay 4,
4 SUNW,emlxs-pci10df,fc20 LPe 11002-S
/pci@3,700000/SUNW,emlxs@0
00 PCIE 3 2, fc20, 10df 2, 0, 1 okay 4,
4 SUNW,emlxs-pci10df,fc20 LPe 11002-S
/pci@3,700000/SUNW,emlxs@0,1
```

- 2 Identify the HBA and its WWPN(s), which you want to replace using the `cfgadm` command.

To identify the HBA, enter the following:

```
# cfgadm -al | grep -i fibre
iou#0-pci#1 fibre/hp connected configured ok
iou#0-pci#4 fibre/hp connected configured ok
```

To list all HBAs, enter the following:

```
# luxadm -e port
/devices/pci@0,600000/pci@0/pci@9/SUNW,emlxs@0/fp@0,0:devctl
NOT CONNECTED
/devices/pci@0,600000/pci@0/pci@9/SUNW,emlxs@0,1/fp@0,0:devctl
CONNECTED
/devices/pci@3,700000/SUNW,emlxs@0/fp@0,0:devctl
NOT CONNECTED
/devices/pci@3,700000/SUNW,emlxs@0,1/fp@0,0:devctl
CONNECTED
```

To select the HBA to dump the portap and get the WWPN, enter the following:

```
# luxadm -e dump_map /devices/pci@0,600000/pci@0/pci@9/SUNW,emlxs@0,1/
fp@0,0:devctl
0      304700    0      203600a0b847900c 200600a0b847900c 0x0
(Disk device)
1      30a800    0      20220002ac00065f 2ff70002ac00065f 0x0
(Disk device)
2      30a900    0      21220002ac00065f 2ff70002ac00065f 0x0
(Disk device)
3      560500    0      10000000c97c3c2f 20000000c97c3c2f 0x1f
(Unknown Type)
4      560700    0      10000000c97c9557 20000000c97c9557 0x1f
(Unknown Type)
5      560b00    0      10000000c97c34b5 20000000c97c34b5 0x1f
(Unknown Type)
6      560900    0      10000000c973149f 20000000c973149f 0x1f
(Unknown Type,Host Bus Adapter)
```

Alternately, you can run the `fcinfo hba-port` Solaris command to get the WWPN(s) for the HBA ports.

- 3 Ensure you have a compatible spare HBA for hot-swap.

- 4 Stop the I/O operations on the HBA port(s) and disable the DMP subpath(s) for the HBA that you want to replace.

```
# vxddmpadm disable ctrl=ctrl1#
```

- 5 Dynamically unconfigure the HBA in the PCIe slot using the `cfgadm` command.

```
# cfgadm -c unconfigure iou#0-pci#1
```

Look for console messages to check if the `cfgadm` command is unsuccessful. If the `cfgadm` command is unsuccessful, proceed to troubleshooting using the server hardware documentation. Check the Solaris 10 patch level recommended for dynamic reconfiguration operations and contact SUN support for further assistance.

```
console messages
```

```
Oct 24 16:21:44 m5000sb0 pcihp: NOTICE: pcihp (pxb_plx2):  
card is removed from the slot iou 0-pci 1
```

- 6 Verify that the HBA card that is being replaced in step 5 is not in the configuration. Enter the following command:

```
# cfgadm -al | grep -i fibre  
iou 0-pci 4 fibre/hp connected configured ok
```

- 7 Mark the fiber cable(s).

- 8 Remove the fiber cable(s) and the HBA that you must replace.

For more information, see the HBA replacement procedures in *SPARC Enterprise M4000/M5000/M8000/M9000 Servers Dynamic Reconfiguration (DR) User's Guide*.

- 9 Replace the HBA with a new compatible HBA of similar type in the same slot. The reinserted card shows up as follows:

```
console messages
```

```
iou 0-pci 1 unknown disconnected unconfigured unknown
```

- 10 Bring the replaced HBA back into the configuration. Enter the following:

```
# cfgadm -c configure iou 0-pci 1
```

```
console messages
```

```
Oct 24 16:21:57 m5000sb0 pcihp: NOTICE: pcihp (pxb_plx2):  
card is inserted in the slot iou#0-pci#1 (pci dev 0)
```

- 11 Verify that the reinserted HBA is in the configuration. Enter the following:

```
# cfgadm -al | grep -i fibre
iou#0-pci 1 fibre/hp connected configured ok <===
iou#0-pci 4 fibre/hp connected configured ok
```

- 12 Modify fabric zoning to include the replaced HBA WWPN(s).
- 13 Enable LUN security on storage for the new WWPN(s).
- 14 Perform an operating system device scan to re-discover the LUNs. Enter the following:

```
# cfgadm -c configure c3
```

- 15 Clean up the device tree for old LUNs. Enter the following:

```
# devfsadm -Cv
```

Note: Sometimes replacing an HBA creates new devices. Perform cleanup operations for the LUN only when new devices are created.

- 16 If SFCFSHA does not show a ghost path for the removed HBA path, enable the path using the `vxddmpadm` command. This performs the device scan for that particular HBA subpath(s). Enter the following:

```
# vxddmpadm enable ctrl=ctrl#
```

- 17 Verify if I/O operations are scheduled on that path. If I/O operations are running correctly on all paths, the dynamic HBA replacement operation is complete.

Upgrading the array controller firmware online

Storage array subsystems need code upgrades as fixes, patches, or feature upgrades. You can perform these upgrades online when the file system is mounted and I/Os are being served to the storage.

Legacy storage subsystems contain two controllers for redundancy. An online upgrade is done one controller at a time. DMP fails over all I/O to the second controller while the first controller is undergoing an Online Controller Upgrade. After the first controller has completely staged the code, it reboots, resets, and comes online with the new version of the code. The second controller goes through the same process, and I/O fails over to the first controller.

Note: Throughout this process, application I/O is not affected.

Array vendors have different names for this process. For example, EMC calls it a nondisruptive upgrade (NDU) for CLARiiON arrays.

A/A type arrays require no special handling during this online upgrade process. For A/P, A/PF, and ALUA type arrays, DMP performs array-specific handling through vendor-specific array policy modules (APMs) during an online controller code upgrade.

When a controller resets and reboots during a code upgrade, DMP detects this state through the SCSI Status. DMP immediately fails over all I/O to the next controller.

If the array does not fully support NDU, all paths to the controllers may be unavailable for I/O for a short period of time. Before beginning the upgrade, set the `dmp_lun_retry_timeout` tunable to a period greater than the time that you expect the controllers to be unavailable for I/O. DMP does not fail the I/Os until the end of the `dmp_lun_retry_timeout` period, or until the I/O succeeds, whichever happens first. Therefore, you can perform the firmware upgrade without interrupting the application I/Os.

For example, if you expect the paths to be unavailable for I/O for 300 seconds, use the following command:

```
# vxddmpadm settune dmp_lun_retry_timeout=300
```

DMP does not fail the I/Os for 300 seconds, or until the I/O succeeds.

To verify which arrays support Online Controller Upgrade or NDU, see the hardware compatibility list (HCL) at the following URL:

<http://www.symantec.com/docs/TECH170013>

Managing devices

This chapter includes the following topics:

- [Displaying disk information](#)
- [Changing the disk device naming scheme](#)
- [About disk installation and formatting](#)
- [Adding and removing disks](#)
- [Renaming a disk](#)

Displaying disk information

Before you use a disk, you need to know if it has been initialized and placed under VxVM control. You also need to know if the disk is part of a disk group, because you cannot create volumes on a disk that is not part of a disk group. The `vxdisk list` command displays device names for all recognized disks, the disk names, the disk group names associated with each disk, and the status of each disk.

To display information on all disks that are known to VxVM

- ◆ Use the following command:

```
# vxdisk list
```

VxVM displays output similar to the following:

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	mydg04	mydg	online
c1t0d0s2	auto:sliced	mydg03	mydg	online
c1t1d0s2	auto:sliced	-	-	online invalid
enc0_2	auto:sliced	mydg02	mydg	online
enc0_3	auto:sliced	mydg05	mydg	online
sena0_0	auto:sliced	-	-	online thinrc1m
sena0_1	auto:sliced	-	-	online

The phrase `online invalid` in the `STATUS` line indicates that a disk has not yet been added to VxVM control. These disks may or may not have been initialized by VxVM previously. Disks that are listed as `online` are already under VxVM control.

To display information about an individual disk

- ◆ Use the following command:

```
# vxdisk [-v] list diskname
```

The `-v` option causes the command to additionally list all tags and tag values that are defined for the disk. By default, tags are not displayed.

Displaying disk information with vxdiskadm

Disk information shows you which disks are initialized, to which disk groups they belong, and the disk status. The `list` option displays device names for all recognized disks, the disk names, the disk group names associated with each disk, and the status of each disk.

To display disk information

- 1 Start the `vxdiskadm` program, and select `list` (List disk information) from the main menu.
- 2 At the following prompt, enter the name of the device you want to see, or enter `all` for a list of all devices:

```
List disk information
Menu: VolumeManager/Disk/ListDisk
```

VxVM INFO V-5-2-475 Use this menu operation to display a list of disks. You can also choose to list detailed information about the disk at a specific disk device address.

Enter disk device or "all" [<address>,all,q,?] (default: all)

- If you enter `all`, VxVM displays the device name, disk name, group, and status of all the devices.
- If you enter the name of a device, VxVM displays complete disk information (including the device name, the type of disk, and information about the public and private areas of the disk) of that device.

Once you have examined this information, press Return to return to the main menu.

Changing the disk device naming scheme

You can either use enclosure-based naming for disks or the operating system's naming scheme. SFCFSA commands display device names according to the current naming scheme.

The default naming scheme is enclosure-based naming (EBN).

When you use DMP with native volumes, the disk naming scheme must be EBN, the `use_avid` attribute must be on, and the persistence attribute must be set to yes.

Note: Devices with very long device names (longer than 31 characters) are represented by enclosure-based names regardless of the naming scheme. If the OS-based names include WWN identifiers, the device name displays with the WWN identifier as long as the device name is less than 31 characters. If any device name is longer than 31 characters, that device name displays with an enclosure name.

To change the disk-naming scheme

- ◆ Select `Change the disk naming scheme` from the `vxdiskadm` main menu to change the disk-naming scheme that you want SFCFSA to use. When prompted, enter `y` to change the naming scheme.

OR

Change the naming scheme from the command line. Use the following command to select enclosure-based naming:

```
# vxddladm set namingscheme=ebn [persistence={yes|no}] \
[use_avid=yes|no] [lowercase=yes|no]
```

Use the following command to select operating system-based naming:

```
# vxddladm set namingscheme=osn [persistence={yes|no}] \
[lowercase=yes|no]
```

The optional `persistence` argument allows you to select whether the names of disk devices that are displayed by SFCFSA remain unchanged after disk hardware has been reconfigured and the system rebooted. By default, enclosure-based naming is persistent. Operating system-based naming is not persistent by default.

To change only the naming persistence without changing the naming scheme, run the `vxddladm set namingscheme` command for the current naming scheme, and specify the persistence attribute.

By default, the names of the enclosure are converted to lowercase, regardless of the case of the name specified by the ASL. The enclosure-based device names are therefore in lower case. Set the `lowercase=no` option to suppress the conversion to lowercase.

For enclosure-based naming, the `use_avid` option specifies whether the Array Volume ID is used for the index number in the device name. By default, `use_avid=yes`, indicating the devices are named as *enclosure_avid*. If `use_avid` is set to `no`, DMP devices are named as *enclosure_index*. The index number is assigned after the devices are sorted by LUN serial number.

The change is immediate whichever method you use.

See [“Regenerating persistent device names”](#) on page 330.

Displaying the disk-naming scheme

SFCFSA disk naming can be operating-system based naming or enclosure-based naming.

The following command displays whether the SFCFSHA disk naming scheme is currently set. It also displays the attributes for the disk naming scheme, such as whether persistence is enabled.

To display the current disk-naming scheme and its mode of operations, use the following command:

```
# vxddladm get namingscheme
```

Setting customized names for DMP nodes

The DMP node name is the meta device name which represents the multiple paths to a disk. The DMP node name is generated from the device name according to the SFCFSHA naming scheme.

You can specify a customized name for a DMP node. User-specified names are persistent even if names persistence is turned off.

You cannot assign a customized name that is already in use by a device. However, if you assign names that follow the same naming conventions as the names that the DDL generates, a name collision can potentially occur when a device is added. If the user-defined name for a DMP device is the same as the DDL-generated name for another DMP device, the `vxdisk list` command output displays one of the devices as 'error'.

To specify a custom name for a DMP node

- ◆ Use the following command:

```
# vxdmpadm setattr dmpnode dmpnodename name=name
```

You can also assign names from an input file. This enables you to customize the DMP nodes on the system with meaningful names.

To assign DMP nodes from a file

- 1 Use the script `vxgetdmpnames` to get a sample file populated from the devices in your configuration. The sample file shows the format required and serves as a template to specify your customized names.
- 2 To assign the names, use the following command:

```
# vxddladm assign names file=pathname
```

To clear custom names

- ◆ To clear the names, and use the default OSN or EBN names, use the following command:

```
# vxddladm -c assign names
```

Regenerating persistent device names

The persistent device naming feature makes the names of disk devices persistent across system reboots. DDL assigns device names according to the persistent device name database.

If operating system-based naming is selected, each disk name is usually set to the name of one of the paths to the disk. After hardware reconfiguration and a subsequent reboot, the operating system may generate different names for the paths to the disks. Therefore, the persistent device names may no longer correspond to the actual paths. This does not prevent the disks from being used, but the association between the disk name and one of its paths is lost.

Similarly, if enclosure-based naming is selected, the device name depends on the name of the enclosure and an index number. If a hardware configuration changes the order of the LUNs exposed by the array, the persistent device name may not reflect the current index.

To regenerate persistent device names

- ◆ To regenerate the persistent names repository, use the following command:

```
# vxddladm [-c] assign names
```

The `-c` option clears all user-specified names and replaces them with autogenerated names.

If the `-c` option is not specified, existing user-specified names are maintained, but OS-based and enclosure-based names are regenerated.

The disk names now correspond to the new path names.

Changing device naming for TPD-controlled enclosures

By default, TPD-controlled enclosures use pseudo device names based on the TPD-assigned node names. If you change the device naming to native, the devices are named in the same format as other SFCFSHA devices. The devices use either operating system names (OSN) or enclosure-based names (EBN), depending on which naming scheme is set.

See [“Displaying the disk-naming scheme”](#) on page 328.

To change device naming for TPD-controlled enclosures

- ◆ For disk enclosures that are controlled by third-party drivers (TPD) whose coexistence is supported by an appropriate ASL, the default behavior is to assign device names that are based on the TPD-assigned node names. You can use the `vxddmpadm` command to switch between these names and the device names that are known to the operating system:

```
# vxddmpadm setattr enclosure enclosure_name tpdmode=native|pseudo
```

The argument to the `tpdmode` attribute selects names that are based on those used by the operating system (`native`), or TPD-assigned node names (`pseudo`).

The use of this command to change between TPD and operating system-based naming is illustrated in the following example for the enclosure named `EMC0`. In this example, the device-naming scheme is set to `OSN`.

```
# vxddisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
emcpower10s2	auto:sliced	disk1	mydg	online
emcpower11s2	auto:sliced	disk2	mydg	online
emcpower12s2	auto:sliced	disk3	mydg	online
emcpower13s2	auto:sliced	disk4	mydg	online
emcpower14s2	auto:sliced	disk5	mydg	online
emcpower15s2	auto:sliced	disk6	mydg	online
emcpower16s2	auto:sliced	disk7	mydg	online
emcpower17s2	auto:sliced	disk8	mydg	online
emcpower18s2	auto:sliced	disk9	mydg	online
emcpower19s2	auto:sliced	disk10	mydg	online

```
# vxddmpadm setattr enclosure EMC0 tpdmode=native
```

```
# vxddisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c6t0d10s2	auto:sliced	disk1	mydg	online
c6t0d11s2	auto:sliced	disk2	mydg	online
c6t0d12s2	auto:sliced	disk3	mydg	online
c6t0d13s2	auto:sliced	disk4	mydg	online
c6t0d14s2	auto:sliced	disk5	mydg	online
c6t0d15s2	auto:sliced	disk6	mydg	online
c6t0d16s2	auto:sliced	disk7	mydg	online
c6t0d17s2	auto:sliced	disk8	mydg	online
c6t0d18s2	auto:sliced	disk9	mydg	online
c6t0d19s2	auto:sliced	disk10	mydg	online

If `tpdmode` is set to `native`, the path with the smallest device number is displayed.

Simple or nopriv disks with enclosure-based naming

If you change from OS-based naming to enclosure-based naming, simple or nopriv disks may be put in the `error` state and cause VxVM objects on those disks to fail.

You can use the `vxdarestore` command to handle simple and nopriv disk failures that arise from changing to the enclosure-based naming scheme. You do not need to use this command if your system does not have any simple or nopriv disks, or if the devices on which any simple or nopriv disks are present are not automatically configured by VxVM (for example, non-standard disk devices such as ramdisks).

Note: You cannot run `vxdarestore` if OS-based naming is in use. Additionally, `vxdarestore` does not handle failures on simple or nopriv disks that are caused by renaming enclosures, by hardware reconfiguration that changes device names, or by changing the naming scheme on a system that includes persistent sliced disk records.

See [“Removing the error state for simple or nopriv disks in the boot disk group”](#) on page 332.

See [“Removing the error state for simple or nopriv disks in non-boot disk groups”](#) on page 333.

See the `vxdarestore(1M)` manual page.

Removing the error state for simple or nopriv disks in the boot disk group

If the boot disk group (usually aliased as `bootdg`) is comprised of only simple and/or nopriv disks, the `vxconfigd` daemon goes into the disabled state after the naming scheme change.

To remove the error state for simple or nopriv disks in the boot disk group

- 1 Use `vxdiskadm` to change back to `c##d##s#` naming.
- 2 Enter the following command to restart the VxVM configuration daemon:

```
# vxconfigd -kr reset
```

- 3 If you want to use enclosure-based naming, use `vxdiskadm` to add a sliced disk to the `bootdg` disk group, change back to the enclosure-based naming scheme, and then run the following command:

```
# vxdarestore
```

Removing the error state for simple or nopriv disks in non-boot disk groups

Removing the error state for simple or nopriv disks in non-boot disk groups

If an imported disk group, other than `bootdg`, is comprised of only simple and/or nopriv disks, the disk group is in the “online `dgdisabled`” state after the change to the enclosure-based naming scheme.

To remove the error state for simple or nopriv disks in non-boot disk groupsTo remove the error state for simple or nopriv disks in disk groups

- 1 Deport the disk group using the following command:
- 2 Use the `vxdarestore` command to restore the failed disks, and to recover the objects on those disks:

```
# vxdarestore
```

- 3 Re-import the disk group using the following command:

```
# vxdg import diskgroup
```

About the Array Volume Identifier (AVID) attribute

DMP assigns enclosure-based names to DMP meta-devices using an array-specific attribute called the Array Volume ID (AVID). The AVID is a unique identifier for the LUN that is provided by the array. The ASL corresponding to the array provides

the AVID property. Within an array enclosure, DMP uses the Array Volume Identifier (AVID) as an index in the DMP metanode name. The DMP metanode name is in the format *enclosureID_AVID*.

The SFCFSA utilities such as `vxdisk list` display the DMP metanode name, which includes the AVID property. Use the AVID to correlate the DMP metanode name to the LUN displayed in the array management interface (GUI or CLI) .

If the ASL does not provide the array volume ID property, then DMP generates an index number. DMP sorts the devices seen from an array by the LUN serial number and then assigns the index number. In this case, the DMP metanode name is in the format *enclosureID_index*.

In a cluster environment, the DMP device names are the same across all nodes in the cluster.

For example, on an EMC CX array where the enclosure is `emc_clariion0` and the array volume ID provided by the ASL is 91, the DMP metanode name is `emc_clariion0_91`. The following sample output shows the DMP metanode names:

```
$ vxdisk list
emc_clariion0_91  auto:cdsdisk  emc_clariion0_91  dg1  online shared
emc_clariion0_92  auto:cdsdisk  emc_clariion0_92  dg1  online shared
emc_clariion0_93  auto:cdsdisk  emc_clariion0_93  dg1  online shared
emc_clariion0_282 auto:cdsdisk  emc_clariion0_282 dg1  online shared
emc_clariion0_283 auto:cdsdisk  emc_clariion0_283 dg1  online shared
emc_clariion0_284 auto:cdsdisk  emc_clariion0_284 dg1  online shared
```

Enclosure based naming with the Array Volume Identifier (AVID) attribute

By default, DMP assigns enclosure-based names to DMP meta-devices using an array-specific attribute called the Array Volume ID (AVID). The AVID provides a unique identifier for the LUN that is provided by the array. The ASL corresponding to the array provides the AVID property. Within an array enclosure, DMP uses the Array Volume Identifier (AVID) as an index in the DMP metanode name. The DMP metanode name is in the format *enclosureID_AVID*.

With the introduction of AVID to the EBN naming scheme, identifying storage devices becomes much easier. The array volume identifier (AVID) enables you to have consistent device naming across multiple nodes connected to the same storage. The disk access name never changes, because it is based on the name defined by the array itself.

Note: DMP does not support AVID with PowerPath names.

If DMP does not have access to a device's AVID, it retrieves another unique LUN identifier called the LUN serial number. DMP sorts the devices based on the LUN Serial Number (LSN), and then assigns the index number. All hosts see the same set of devices, so all hosts will have the same sorted list, leading to consistent device indices across the cluster. In this case, the DMP metanode name is in the format `enclosureID_index`.

DMP also supports a scalable framework, that allows you to fully customize the device names on a host by applying a device naming file that associates custom names with cabinet and LUN serial numbers.

If a CVM cluster is symmetric, each node in the cluster accesses the same set of disks. Enclosure-based names provide a consistent naming system so that the device names are the same on each node.

The SFCFSA utilities such as `vxdisk list` display the DMP metanode name, which includes the AVID property. Use the AVID to correlate the DMP metanode name to the LUN displayed in the array management interface (GUI or CLI) .

For example, on an EMC CX array where the enclosure is `emc_clariion0` and the array volume ID provided by the ASL is 91, the DMP metanode name is `emc_clariion0_91`. The following sample output shows the DMP metanode names:

```
$ vxdisk list
emc_clariion0_91 auto:cdsdisk emc_clariion0_91 dg1 online shared
emc_clariion0_92 auto:cdsdisk emc_clariion0_92 dg1 online shared
emc_clariion0_93 auto:cdsdisk emc_clariion0_93 dg1 online shared
emc_clariion0_282 auto:cdsdisk emc_clariion0_282 dg1 online shared
emc_clariion0_283 auto:cdsdisk emc_clariion0_283 dg1 online shared
emc_clariion0_284 auto:cdsdisk emc_clariion0_284 dg1 online shared

# vxddladm get namingscheme
NAMING_SCHEME      PERSISTENCE    LOWERCASE      USE_AVID
=====
Enclosure Based    Yes             Yes             Yes
```

About disk installation and formatting

Depending on the hardware capabilities of your disks and of your system, you may either need to shut down and power off your system before installing the disks, or you may be able to hot-insert the disks into the live system. Many operating systems can detect the presence of the new disks on being rebooted. If the disks are inserted while the system is live, you may need to enter an operating system-specific command to notify the system.

If the disks require low or intermediate-level formatting before use, use the operating system-specific formatting command to do this.

For Solaris SPARC systems, if a disk is not formatted, the status field in the `vxdisk list` output shows as `nolabel`. The disk must be formatted before you add it to VxVM control.

Note: SCSI disks are usually preformatted. Reformatting is needed only if the existing formatting has become damaged.

See [“Adding a disk to VxVM”](#) on page 336.

Adding and removing disks

This section describes managing devices.

Adding a disk to VxVM

Formatted disks being placed under VxVM control may be new or previously used outside VxVM. The set of disks can consist of all disks on the system, all disks on a controller, selected disks, or a combination of these.

Depending on the circumstances, all of the disks may not be processed in the same way.

For example, some disks may be initialized, while others may be encapsulated to preserve existing data on the disks.

When initializing multiple disks at one time, it is possible to exclude certain disks or certain controllers.

You can also exclude certain disks or certain controllers when encapsulating multiple disks at one time.

To exclude a device from the view of VxVM, select `Prevent multipathing/Suppress devices` from VxVM's view from the `vxdiskadm` main menu.

Warning: Initialization does not preserve the existing data on the disks.

See [“Making devices invisible to VxVM”](#) on page 262.

To initialize disks for VxVM use

- 1 Select Add or initialize one or more disks from the `vxdiskadm` main menu.
- 2 At the following prompt, enter the disk device name of the disk to be added to VxVM control (or enter `list` for a list of disks):

```
Select disk devices to add:
[<pattern-list>,all,list,q,?]
```

The *pattern-list* can be a single disk, or a series of disks and/or controllers (with optional targets). If *pattern-list* consists of multiple items, separate them using white space. For example, specify four disks at separate target IDs on controller 3 as follows:

```
c3t0d0 c3t1d0 c3t2d0 c3t3d0
```

If you enter `list` at the prompt, the `vxdiskadm` program displays a list of the disks available to the system:

DEVICE	DISK	GROUP	STATUS
c0t0d0	mydg01	mydg	online
c0t1d0	mydg02	mydg	online
c0t2d0	mydg03	mydg	online
c0t3d0	-	-	online
c1t0d0	mydg10	mydg	online
c1t0d1	-	-	online invalid
.			
.			
.			
c3t0d0	-	-	online invalid
sena0_0	mydg33	mydg	online
sena0_1	mydg34	mydg	online
sena0_2	mydg35	mydg	online

The phrase `online invalid` in the `STATUS` line indicates that a disk has yet to be added or initialized for VxVM control. Disks that are listed as `online` with a disk name and disk group are already under VxVM control.

Enter the device name or pattern of the disks that you want to initialize at the prompt and press Return.

- 3 To continue with the operation, enter `y` (or press Return) at the following prompt:

```
Here are the disks selected. Output format: [Device]
list of device names
```

```
Continue operation? [y,n,q,?] (default: y) y
```

- 4 At the following prompt, specify the disk group to which the disk should be added, or `none` to reserve the disks for future use:

```
You can choose to add these disks to an existing disk group,
a new disk group, or you can leave these disks available for use
by future add or replacement operations. To create a new disk
group, select a disk group name that does not yet exist. To
leave the disks available for future use, specify a disk group
name of none.
```

```
Which disk group [<group>,none,list,q,?]
```

- 5 If you specified the name of a disk group that does not already exist, `vxdiskadm` prompts for confirmation that you really want to create this new disk group:

```
There is no active disk group named disk group name.
```

```
Create a new group named disk group name? [y,n,q,?]
(default: y) y
```

You are then prompted to confirm whether the disk group should support the Cross-platform Data Sharing (CDS) feature:

```
Create the disk group as a CDS disk group? [y,n,q,?]
(default: n)
```

If the new disk group may be moved between different operating system platforms, enter `y`. Otherwise, enter `n`.

- 6 At the following prompt, either press Return to accept the default disk name or enter `n` to allow you to define your own disk names:

```
Use default disk names for the disks? [y,n,q,?] (default: y) n
```

- 7 When prompted whether the disks should become hot-relocation spares, enter `n` (or press Return):

```
Add disks as spare disks for disk group name? [y,n,q,?]
(default: n) n
```

- 8 When prompted whether to exclude the disks from hot-relocation use, enter **n** (or press Return).

```
Exclude disks from hot-relocation use? [y,n,q,?]
(default: n) n
```

- 9 You are next prompted to choose whether you want to add a site tag to the disks:

```
Add site tag to disks? [y,n,q,?] (default: n)
```

A site tag is usually applied to disk arrays or enclosures, and is not required unless you want to use the Remote Mirror feature.

If you enter **y** to choose to add a site tag, you are prompted to the site name at step 11.

- 10 To continue with the operation, enter **y** (or press Return) at the following prompt:

```
The selected disks will be added to the disk group
disk group name with default disk names.
list of device names
Continue with operation? [y,n,q,?] (default: y) y
```

- 11 If you chose to tag the disks with a site in step 9, you are now prompted to enter the site name that should be applied to the disks in each enclosure:

```
The following disk(s):
list of device names

belong to enclosure(s):
list of enclosure names

Enter site tag for disks on enclosure enclosure name
[<name>,q,?] site_name
```

- 12** If you see the following prompt, it lists any disks that have already been initialized for use by VxVM:

The following disk devices appear to have been initialized already.

The disks are currently available as replacement disks.

Output format: [Device]

list of device names

Use these devices? [Y,N,S(select),q,?] (default: Y) **Y**

This prompt allows you to indicate “yes” or “no” for all of these disks (Y or N) or to select how to process each of these disks on an individual basis (S).

If you are sure that you want to reinitialize all of these disks, enter Y at the following prompt:

VxVM NOTICE V-5-2-366 The following disks you selected for use appear to already have been initialized for the Volume Manager. If you are certain the disks already have been initialized for the Volume Manager, then you do not need to reinitialize these disk devices.

Output format: [Device]

list of device names

Reinitialize these devices? [Y,N,S(select),q,?] (default: Y) **Y**

- 13** `vxdiskadm` may now indicate that one or more disks is a candidate for encapsulation. Encapsulation allows you to add an active disk to VxVM control and preserve the data on that disk. If you want to preserve the data on the disk, enter `y`. If you are sure that there is no data on the disk that you want to preserve, enter `n` to avoid encapsulation.

```
VxVM NOTICE V-5-2-355 The following disk device has a valid
partition table, but does not appear to have been initialized
for the Volume Manager. If there is data on the disk that
should NOT be destroyed you should encapsulate the existing
disk partitions as volumes instead of adding the disk as a new
disk.
```

```
Output format: [Device]
```

```
device name
```

```
Encapsulate this device? [y,n,q,?] (default: y)
```

- 14 If you choose to encapsulate the disk `vxdiskadm` confirms its device name and prompts you for permission to proceed. Enter `y` (or press Return) to continue encapsulation:

```
VxVM NOTICE V-5-2-311 The following disk device has been
selected for encapsulation.
```

```
Output format: [Device]
```

```
device name
```

```
Continue with encapsulation? [y,n,q,?] (default: y) y
```

```
vxdiskadm now displays an encapsulation status and informs you
that you must perform a shutdown and reboot as soon as
possible:
```

```
VxVM INFO V-5-2-333 The disk device device name will be
encapsulated and added to the disk group disk group name with the
disk name disk name.
```

You can now choose whether the disk is to be formatted as a CDS disk that is portable between different operating systems, or as a non-portable sliced or simple disk:

```
Enter the desired format [cdsdisk,sliced,simple,q,?]
(default: cdsdisk)
```

Enter the format that is appropriate for your needs. In most cases, this is the default format, `cdsdisk`.

At the following prompt, `vxdiskadm` asks if you want to use the default private region size of 65536 blocks (32MB). Press Return to confirm that you want to use the default value, or enter a different value. (The maximum value that you can specify is 524288 blocks.)

```
Enter desired private region length [<privlen>,q,?]
(default: 65536)
```

If you entered `cdsdisk` as the format, you are prompted for the action to be taken if the disk cannot be converted this format:

```
Do you want to use sliced as the format should cdsdisk fail?
[y,n,q,?] (default: y)
```

If you enter `y`, and it is not possible to encapsulate the disk as a CDS disk, it is encapsulated as a sliced disk. Otherwise, the encapsulation fails.

`vxdiskadm` then proceeds to encapsulate the disks. You should now reboot your system at the earliest possible opportunity, for example by running this command:

```
# shutdown -g0 -y -i6
```

The `/etc/vfstab` file is updated to include the volume devices that are used to mount any encapsulated file systems. You may need to update any other references in backup scripts, databases, or manually created swap devices. The original `/etc/vfstab` file is saved as `/etc/vfstab.prevm`.

- 15** If you choose not to encapsulate the disk `vxdiskadm` asks if you want to initialize the disk instead. Enter `y` to confirm this:

Instead of encapsulating, initialize? [y,n,q,?] (default: n) `y` `vxdiskadm` now confirms those disks that are being initialized and added to VxVM control with messages similar to the following. In addition, you may be prompted to perform surface analysis.

```
VxVM INFO V-5-2-205 Initializing device device name.
```

- 16** You can now choose whether the disk is to be formatted as a CDS disk that is portable between different operating systems, or as a non-portable sliced or simple disk:

```
Enter the desired format [cdsdisk,sliced,simple,q,?]
(default: cdsdisk)
```

Enter the format that is appropriate for your needs. In most cases, this is the default format, `cdsdisk`.

- 17** At the following prompt, `vxdiskadm` asks if you want to use the default private region size of 65536 blocks (32MB). Press Return to confirm that you want to use the default value, or enter a different value. (The maximum value that you can specify is 524288 blocks.)

`vxdiskadm` then proceeds to add the disks.

```
VxVM INFO V-5-2-88 Adding disk device name to disk group
disk group name with disk name disk name.
```

```
.
.
.
```

- 18 If you choose not to use the default disk names, `vxdiskadm` prompts you to enter the disk name.
- 19 At the following prompt, indicate whether you want to continue to initialize more disks (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Add or initialize other disks? [y,n,q,?] (default: n)
```

The default layout for disks can be changed.

Disk reinitialization

You can reinitialize a disk that has previously been initialized for use by VxVM by putting it under VxVM control as you would a new disk.

See [“Adding a disk to VxVM”](#) on page 336.

Warning: Reinitialization does not preserve data on the disk. If you want to reinitialize the disk, make sure that it does not contain data that should be preserved.

If the disk you want to add has been used before, but not with VxVM, you can encapsulate the disk to preserve its information. If the disk you want to add has previously been under the control of Solaris Volume Manager, you can preserve the data it contains on a VxVM disk by the process of conversion.

For detailed information about migrating volumes, see the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.

Using `vxdiskadd` to put a disk under VxVM control

To use the `vxdiskadd` command to put a disk under VxVM control.

- ◆ Type the following command:

```
# vxdiskadd disk
```

For example, to initialize the second disk on the first controller:

```
# vxdiskadd c0t1d0
```

The `vxdiskadd` command examines your disk to determine whether it has been initialized and also checks for disks that have been added to VxVM, and for other conditions.

The `vxdiskadd` command also checks for disks that can be encapsulated.

See [“Encapsulating a disk”](#) on page 866.

If you are adding an uninitialized disk, warning and error messages are displayed on the console by the `vxdiskadd` command. Ignore these messages. These messages should not appear after the disk has been fully initialized; the `vxdiskadd` command displays a success message when the initialization completes.

The interactive dialog for adding a disk using `vxdiskadd` is similar to that for `vxdiskadm`.

See [“Adding a disk to VxVM”](#) on page 336.

Removing disks

You must disable a disk group before you can remove the last disk in that group.

As an alternative to disabling the disk group, you can destroy the disk group.

You can remove a disk from a system and move it to another system if the disk is failing or has failed.

To remove a disk

- 1 Stop all activity by applications to volumes that are configured on the disk that is to be removed. Unmount file systems and shut down databases that are configured on the volumes.
- 2 Use the following command to stop the volumes:

```
# vxvol [-g diskgroup] stop vol1 vol2 ...
```

- 3 Move the volumes to other disks or back up the volumes. To move a volume, use `vxdiskadm` to mirror the volume on one or more disks, then remove the original copy of the volume. If the volumes are no longer needed, they can be removed instead of moved.
- 4 Check that any data on the disk has either been moved to other disks or is no longer needed.
- 5 Select `Remove a disk` from the `vxdiskadm` main menu.
- 6 At the following prompt, enter the disk name of the disk to be removed:

```
Enter disk name [<disk>,list,q,?] mydg01
```

- 7 If there are any volumes on the disk, VxVM asks you whether they should be evacuated from the disk. If you wish to keep the volumes, answer `y`. Otherwise, answer `n`.
- 8 At the following verification prompt, press Return to continue:

```
VxVM NOTICE V-5-2-284 Requested operation is to remove disk  
mydg01 from group mydg.
```

```
Continue with operation? [y,n,q,?] (default: y)
```

The `vxdiskadm` utility removes the disk from the disk group and displays the following success message:

```
VxVM INFO V-5-2-268 Removal of disk mydg01 is complete.
```

You can now remove the disk or leave it on your system as a replacement.

- 9 At the following prompt, indicate whether you want to remove other disks (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Remove another disk? [y,n,q,?] (default: n)
```

Removing a disk with subdisks

You can remove a disk on which some subdisks are defined. For example, you can consolidate all the volumes onto one disk. If you use the `vxdiskadm` program to remove a disk, you can choose to move volumes off that disk.

Some subdisks are not movable. A subdisk may not be movable for one of the following reasons:

- There is not enough space on the remaining disks in the subdisks disk group.

- Plexes or striped subdisks cannot be allocated on different disks from existing plexes or striped subdisks in the volume.

If the `vxdiskadm` program cannot move some subdisks, remove some plexes from some disks to free more space before proceeding with the disk removal operation.

See [“Removing a volume”](#) on page 854.

To remove a disk with subdisks

- 1 Run the `vxdiskadm` program and select `Remove a disk` from the main menu.

If the disk is used by some subdisks, the following message is displayed:

```
VxVM ERROR V-5-2-369 The following volumes currently use part of
disk mydg02:
```

```
home usrvol
```

```
Volumes must be moved from mydg02 before it can be removed.
```

```
Move volumes to other disks? [y,n,q,?] (default: n)
```

- 2 Choose `y` to move all subdisks off the disk, if possible.

Removing a disk with no subdisks

To remove a disk that contains no subdisks from its disk group

- ◆ Run the `vxdiskadm` program and select `Remove a disk` from the main menu, and respond to the prompts as shown in this example to remove `mydg02`:

```
Enter disk name [<disk>,list,q,?] mydg02
```

```
VxVM NOTICE V-5-2-284 Requested operation is to remove disk
mydg02 from group mydg.
```

```
Continue with operation? [y,n,q,?] (default: y) y
```

```
VxVM INFO V-5-2-268 Removal of disk mydg02 is complete.
```

```
Clobber disk headers? [y,n,q,?] (default: n) y
```

Enter `y` to remove the disk completely from VxVM control. If you do not want to remove the disk completely from VxVM control, enter `n`.

Renaming a disk

If you do not specify a VM disk name, VxVM gives the disk a default name when you add the disk to VxVM control. The VM disk name is used by VxVM to identify the location of the disk or the disk type.

To rename a disk

- ◆ Type the following command:

```
# vxedit [-g diskgroup] rename old_diskname new_diskname
```

By default, VxVM names subdisk objects after the VM disk on which they are located. Renaming a VM disk does not automatically rename the subdisks on that disk.

For example, you might want to rename disk `mydg03`, as shown in the following output from `vxdisk list`, to `mydg02`:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	mydg01	mydg	online
c1t0d0s2	auto:sliced	mydg03	mydg	online
c1t1d0s2	auto:sliced	-	-	online

You would use the following command to rename the disk.

```
# vxedit -g mydg rename mydg03 mydg02
```

To confirm that the name change took place, use the `vxdisk list` command again:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:sliced	mydg01	mydg	online
c1t0d0s2	auto:sliced	mydg02	mydg	online
c1t1d0s2	auto:sliced	-	-	online

Event monitoring

This chapter includes the following topics:

- [About the Dynamic Multi-Pathing \(DMP\) event source daemon \(vxesd\)](#)
- [Fabric Monitoring and proactive error detection](#)
- [Dynamic Multi-Pathing \(DMP\) automated device discovery on Solaris](#)
- [Dynamic Multi-Pathing \(DMP\) discovery of iSCSI and SAN Fibre Channel topology](#)
- [DMP event logging](#)
- [Starting and stopping the Dynamic Multi-Pathing \(DMP\) event source daemon](#)

About the Dynamic Multi-Pathing (DMP) event source daemon (vxesd)

The event source daemon (`vxesd`) is a Veritas Dynamic Multi-Pathing (DMP) component process that receives notifications of any device-related events that are used to take appropriate actions. The benefits of `vxesd` include:

- Monitoring of SAN fabric events and proactive error detection (SAN event)
See [“Fabric Monitoring and proactive error detection”](#) on page 350.
- Logging of DMP events for troubleshooting (DMP event)
See [“DMP event logging”](#) on page 352.
- On Solaris:
Automated device discovery (OS event)
See [“Dynamic Multi-Pathing \(DMP\) automated device discovery on Solaris”](#) on page 351.

- Discovery of SAN components and HBA-array port connectivity (Fibre Channel and iSCSI)
See [“Dynamic Multi-Pathing \(DMP\) discovery of iSCSI and SAN Fibre Channel topology”](#) on page 352.
See [“Starting and stopping the Dynamic Multi-Pathing \(DMP\) event source daemon”](#) on page 353.

Fabric Monitoring and proactive error detection

In previous releases, DMP handled failed paths reactively, by only disabling paths when active I/O failed on the storage. Using the Storage Networking Industry Association (SNIA) HBA API library, `vxesd` now is able to receive SAN fabric events from the HBA. This information allows DMP to take a proactive role by checking suspect devices from the SAN events, even if there is no active I/O. New I/O is directed to healthy paths while the suspect devices are verified.

During startup, `vxesd` queries the HBA (by way of the SNIA library) to obtain the SAN topology. The `vxesd` daemon determines the Port World Wide Names (PWWN) that correspond to each of the device paths that are visible to the operating system. After the `vxesd` daemon obtains the topology, `vxesd` registers with the HBA for SAN event notification. If LUNs are disconnected from a SAN, the HBA notifies `vxesd` of the SAN event, specifying the PWWNs that are affected. The `vxesd` daemon uses this event information and correlates it with the previous topology information to determine which set of device paths have been affected.

The `vxesd` daemon sends the affected set to the `vxconfigd` daemon (DDL) so that the device paths can be marked as suspect. When the path is marked as suspect, DMP does not send new I/O to the path unless it is the last path to the device. In the background, the DMP restore task checks the accessibility of the paths on its next periodic cycle using a SCSI inquiry probe. If the SCSI inquiry fails, DMP disables the path to the affected LUNs, which is also logged in the event log.

If the LUNs are reconnected at a later time, the HBA informs `vxesd` of the SAN event. When the DMP restore task runs its next test cycle, the disabled paths are checked with the SCSI probe and re-enabled if successful.

Note: If `vxesd` receives an HBA LINK UP event, the DMP restore task is restarted and the SCSI probes run immediately, without waiting for the next periodic cycle. When the DMP restore task is restarted, it starts a new periodic cycle. If the disabled paths are not accessible by the time of the first SCSI probe, they are re-tested on the next cycle (300s by default).

The fabric monitor functionality is enabled by default. The value of the `dmp_monitor_fabric` tunable is persistent across reboots.

To disable the Fabric Monitoring functionality, use the following command:

```
# vxddmpadm settune dmp_monitor_fabric=off
```

To enable the Fabric Monitoring functionality, use the following command:

```
# vxddmpadm settune dmp_monitor_fabric=on
```

To display the current value of the `dmp_monitor_fabric` tunable, use the following command:

```
# vxddmpadm gettune dmp_monitor_fabric
```

Dynamic Multi-Pathing (DMP) automated device discovery on Solaris

In releases before VxVM 4.0, VxVM device discovery required manual invocation of commands such as `vxdisk scandisks` or `vxctl enable.vxesd` automates the discovery process by interfacing with the Reconfiguration Coordination Manager (RCM) framework.

The `vxesd` daemon registers the script `es_devfs.pl` with the Solaris `syseventd` daemon for device arrival events. In the event that `cfgadm` is invoked to attach a new device to the system, the `syseventd` daemon executes the scripts that are registered for device arrival events, including `es_devfs.pl`. The `es_devfs.pl` script establishes a socket with `vxesd` and transfers the event parameter (physical path of device) to the daemon. The `vxesd` daemon in turn connects to the `vxconfigd` daemon to initiate DDL device discovery for the device that had arrived. The whole operation takes place asynchronously so that the `cfgadm` command returns after the event has been added to the `syseventd` queue.

In the event that a device is removed with `cfgadm`, a similar process exists which uses the `es_rcm.pl` script to disable the relevant DMP paths. The removal operation is synchronous so that the `cfgadm` command waits until all the registered detach scripts have completed execution.

Note: On systems with EMC PowerPath, a slow PowerPath discovery process may lead to a device being automatically claimed and controlled by DMP control. In such scenarios, the `vxesd` daemon may be stopped before the addition of the disk and restart after PowerPath has claimed control of the device.

Dynamic Multi-Pathing (DMP) discovery of iSCSI and SAN Fibre Channel topology

The `vxesd` builds a topology of iSCSI and Fibre Channel devices that are visible to the host.

The `vxesd` daemon uses the iSCSI management API (IMA) to build the topology.

To display the hierarchical listing of Fibre Channel and iSCSI devices, use the following command:

```
# vxddladm list
```

See the `vxddladm(1M)` manual page.

DMP event logging

DMP notifies `vxesd` of major events, and `vxesd` logs the event in a log file (`/etc/vx/dmpevents.log`). These events include:

- Marking paths or dmpnodes enabled
- Marking paths or dmpnodes disabled
- Throttling of paths
- I/O error analysis
- Host Bus Adapter (HBA) and Storage Area Network (SAN) events

The log file is located in `/var/adm/vx/dmpevents.log` but is symbolically linked to `/etc/vx/dmpevents.log`. When the file reaches 10,000 lines, the log is rotated.

That is, `dmpevents.log` is renamed `dmpevents.log.X` and a new `dmpevents.log` file is created.

You can change the level of detail in the event log file using the tunable `dmp_log_level`. Valid values are 1 through 4. The default level is 1.

```
# vxdmpadm settune dmp_log_level=X
```

The current value of `dmp_log_level` can be displayed with:

```
# vxdmpadm gettune dmp_log_level
```

For details on the various log levels, see the `vxdmpadm(1M)` manual page.

Starting and stopping the Dynamic Multi-Pathing (DMP) event source daemon

By default, DMP starts `vxesd` at boot time.

To stop the `vxesd` daemon, use the `vxddladm` utility:

```
# vxddladm stop eventsource
```

To start the `vxesd` daemon, use the `vxddladm` utility:

```
# vxddladm start eventsource [logfile=logfilename]
```

To view the status of the `vxesd` daemon, use the `vxddladm` utility:

```
# vxddladm status eventsource
```

Administering Storage Foundation Cluster File System High Availability and its components

- [Chapter 15. Administering Storage Foundation Cluster File System High Availability and its components](#)
- [Chapter 16. Administering sites and remote mirrors](#)

Administering Storage Foundation Cluster File System High Availability and its components

This chapter includes the following topics:

- [About Storage Foundation Cluster File System High Availability administration](#)
- [Administering CFS](#)
- [Administering VCS](#)
- [Administering CVM](#)
- [Administering ODM](#)
- [Administering I/O Fencing](#)
- [Administering SFCFSHA global clusters](#)

About Storage Foundation Cluster File System High Availability administration

The Storage Foundation Cluster File System High Availability (SFCFSHA) is a shared file system that enables multiple hosts to mount and perform file operations concurrently on the same file. To operate in a cluster configuration, SFCFSHA requires the integrated set of Veritas products included in the Veritas Storage Foundation Cluster File System High Availability.

To configure a cluster, SFCFSHA requires the Veritas Cluster Server (VCS). VCS supplies two major components integral to SFCFSHA. The LLT package provides node-to-node communications and monitors network communications. The GAB package provides cluster state, configuration, and membership service, and monitors the heartbeat links between systems to ensure that they are active. There are several other packages supplied by VCS that provide application failover support when installing SFCFSHA.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

SFCFSHA also requires the cluster functionality (CVM) of the Veritas Volume Manager (VxVM) to create the shared volumes necessary for mounting cluster file systems.

For more information on these products, refer to the *Veritas Volume Manager* and *Veritas Cluster Server* documentation.

Administering CFS

This section describes some of the major aspects of cluster file system administration.

This section provides instructions for the following CFS administration tasks:

- Adding CFS file systems to VCS configuration
See [“Adding CFS file systems to VCS configuration”](#) on page 356.
- Resizing CFS file systems
See [“Resizing CFS file systems”](#) on page 357.
- Verifying the status of CFS file systems
See [“Verifying the status of CFS file systems”](#) on page 357.

If you encounter issues while administering CFS, refer to the troubleshooting section for assistance.

Adding CFS file systems to VCS configuration

Run the following command to add a CFS file system to the VCS main.cf file without using an editor.

For example:

```
# cfsmntadm add oradatadg oradatavol \  
/oradata1 all=suid,rw
```



```
Mount Point is being added...
/oradata1 added to the cluster-configuration
```

Using cfsmount to mount CFS file systems

To mount a CFS file system using cfsmount:

```
# cfsmntadm add sdg vol1 /oradata1 all=
# cfsmount /oradata1
Mounting...
[/dev/vx/dsk/oradatadg/oradatavol]
mounted successfully at /oradata1 on system01
[/dev/vx/dsk/oradatadg/oradatavol]
mounted successfully at /oradata1 on system02
```

Resizing CFS file systems

If you see a message on the console indicating that a CFS file system is full, you may want to resize the file system. The `vxresize` command lets you resize a CFS file system. It extends the file system and the underlying volume.

See the `vxresize` (1M) manual page for information on various options.

The following command resizes an Oracle data CFS file system (the Oracle data volume is CFS mounted):

```
# vxresize -g oradatadg oradatavol +2G
```

Verifying the status of CFS file systems

Run the "cfscluster status" command to see the status of the nodes and their mount points:

```
# cfscluster status
```

```
Node           : system02
Cluster Manager : not-running
CVM state       : not-running

MOUNT POINT    SHARED VOLUME  DISK GROUP      STATUS
-----
/ocrvote       ocrvotevol     system01_ocr     NOT MOUNTED
/oracle        ora_vol        system01_ora     NOT MOUNTED
/crshome       ora_crs_vol    system01_crs     NOT MOUNTED
/oradata1      ora_data1_vol  system01_data1   NOT MOUNTED
/arch          archivol       system01_data1   NOT MOUNTED
```

```

Node                : system01
Cluster Manager    : running
CVM state          : running
MOUNT POINT        SHARED VOLUME  DISK GROUP          STATUS
/ocrvote            ocrvotevol     system01_ocr         MOUNTED
/oracle             ora_vol       system01_ora         MOUNTED
/crshome            ora_crs_vol   system01_crs         MOUNTED
/oradata1           ora_data1_vol system01_data1        MOUNTED
/arch               archivol     system01_data1        MOUNTED

```

Verifying CFS port

CFS uses port f for communication between nodes. Port f is the GAB port used for CFS membership.

The CFS port state can be verified as follows:

```
# gabconfig -a | grep "Port f"
```

CFS agents and AMF support

The CFS agents (CFSMount and CFSfsckd) are AMF-aware. In this release the CFS agents use the V51 framework.

CFS agent log files

You can use the CFS agent log files that are located in the directory /var/VRTSvcS/log to debug CFS issues.

```
# cd /var/VRTSvcS/log
# ls
CFSMount_A.log
CFSfsckd_A.log
engine_A.log

```

The agent framework information is located in the engine_A.log file while the agent entry point information is located in the CFSMount_A.log and CFSfsckd_A.log files.

Storage Foundation Cluster File System High Availability commands

[Table 15-1](#) describes the SFCFSHA commands.

Table 15-1 SFCFSA commands

Commands	Description
<code>cfsccluster</code>	Cluster configuration command See the <code>cfsccluster(1M)</code> manual page for more information.
<code>cfsmntadm</code>	Adds, deletes, modifies, and sets policy on cluster mounted file systems See the <code>cfsmntadm(1M)</code> manual page for more information.
<code>cfsgdadm</code>	Adds or deletes shared disk groups to and from a cluster configuration See the <code>cfsgdadm(1M)</code> manual page for more information.
<code>cfsmount</code>	Mounts a cluster file system on a shared volume See the <code>cfsmount(1M)</code> manual page for more information.
<code>cfsumount</code>	Unmounts a cluster file system on a shared volume See the <code>cfsumount(1M)</code> manual page for more information.
<code>cfsshare</code>	Clustered NFS (CNFS) and Common Internet File System (CIFS) configuration command See the <code>cfsshare(1M)</code> manual page for more information.

mount, fsclustadm, and fsadm commands

The `fsadm` command performs online administration functions on VxFS file systems, Storage Checkpoints, or individual files and directories. The `fsadm` command supports file system resizing, extent reorganization, directory reorganization, querying or changing the `largefiles` flag, and Thin Storage Reclamation. The `fsadm` command operates on file systems mounted for read/write access; however, the `-o` option can also operate on a special device containing a clean, unmounted file system. Only a privileged user can change the `largefiles` flag on a mounted file system, or resize or reorganize a file system. You can invoke only one instance of `fsadm` per file system at a time. If `mount_point` is a Storage Checkpoint, `fsadm` performs the specified operation on the entire file system, including all of the file system's Storage Checkpoints.

The `mount` and `fsclustadm` commands are important for configuring cluster file systems.

mount

The `mount` command with the `-o cluster` option lets you access shared file systems.

See the `mount_vxfs(1M)` manual page.

fsclustadm

The `fsclustadm` command reports various attributes of a cluster file system. Using `fsclustadm` you can show and set the primary node in a cluster, translate node IDs to host names and vice versa, list all nodes that currently have a cluster mount of the specified file system mount point, and determine whether a mount is a local or cluster mount. The `fsclustadm` command operates from any node in a cluster on which the file system is mounted, and can control the location of the primary for a specified mount point.

See the `fsclustadm(1M)` manual page.

fsadm

The `fsadm` command performs online administration functions on VxFS file systems, Storage Checkpoints, or individual files and directories. The `fsadm` command supports file system resizing, extent reorganization, directory reorganization, querying or changing the `largefiles` flag, Thin Storage Reclamation, and free space defragmentation.

The `fsadm` command can be invoked from the primary or secondary node.

See the `fsadm_vxfs(1M)` manual page.

Run commands safely in a cluster environment

Any UNIX command that can write to a raw device must be used carefully in a shared environment to prevent data from being corrupted. For shared VxVM volumes, SFCFSA provides protection by reserving the volumes in a cluster to prevent VxFS commands, such as `fsck` and `mkfs`, from inadvertently damaging a mounted file system from another node in a cluster. However, commands such as `dd` execute without any reservation, and can damage a file system mounted from another node. Before running this kind of command on a file system, be sure the file system is not mounted on a cluster. You can run the `mount` command to see if a file system is a shared or local mount.

Time synchronization for Cluster File Systems

SFCFSA requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for creation (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

Growing a Storage Foundation Cluster File System High Availability

There is a master node for CVM as well as a primary for SFCFSHA. When growing a file system, you grow the volume and slaves from the CVM master, and then grow the file system from any SFCFSHA node. The CVM master and the SFCFSHA node can be different nodes.

To determine the primary file system in a cluster

- To determine the primary file system in a cluster, type the following command:

```
# fsclustadm -v showprimary mount_point
```

To determine that the current node is the master CVM node

- To determine if the current node is the master CVM node, type the following command:

```
# vxdctl -c mode
```

To actually increase the size of the file system

- 1 On the master CVM node, type the following command:

```
# vxassist -g shared_disk_group growto volume_name newlength
```

- 2 On any SFCFSHA node, type the following command:

```
# fsadm -F vxfs -b newsize -r device_name mount_point
```

The /etc/vfstab file

In the `/etc/vfstab` file, do not specify any cluster file systems to mount-at-boot because mounts initiated from `vfstab` occur before cluster configuration begins. For cluster mounts, use the VCS configuration file to determine which file systems to enable following a reboot.

When the Storage Foundation Cluster File System High Availability primary fails

If the server on which the SFCFSHA primary is running fails, the remaining cluster nodes elect a new primary. The new primary reads the file system intent log and completes any metadata updates that were in process at the time of the failure. Application I/O from other nodes may block during this process and cause a delay. When the file system is again consistent, application processing resumes.

Because nodes using a cluster file system in secondary node do not update file system metadata directly, failure of a secondary node does not require metadata repair. SFCFSHA recovery from secondary node failure is therefore faster than from primary node failure.

See the `fsclustadm(1M)` manual page for more information.

Distribute the load on a cluster

Distributing the workload in a cluster provides performance and failover advantages.

For example, if you have eight file systems and four nodes, designating two file systems per node as the primary would be beneficial. Primaryship is determined by which node first mounts the file system. You can also use the `fsclustadm` to designate a SFCFSHA primary. The `fsclustadm setprimary` command can also define the order in which primaryship is assumed if the current primary fails. After setup, the policy is in effect as long as one or more nodes in the cluster have the file system mounted.

Storage Checkpoints on Storage Foundation Cluster File System High Availability

The creation of Storage Checkpoints works the same on CFS file systems as they do on local mount file systems.

More information on how to create and maintain Storage Checkpoints is available.

See [“About Storage Checkpoints”](#) on page 640.

Snapshots on Storage Foundation Cluster File System High Availability

A snapshot provides a consistent point-in-time image of a VxFS file system. A snapshot can be accessed as a read-only mounted file system to perform efficient online backups of the file system. Snapshots implement copy-on-write semantics that incrementally copy data blocks when they are overwritten on the snapped file system.

See [“About snapshot file systems”](#) on page 583.

Snapshots for cluster file systems extend the same copy-on-write mechanism for the I/O originating from any node in the cluster.

Cluster snapshot characteristics

A cluster snapshot has the following characteristics:

- A snapshot for a cluster mounted file system can be mounted on any node in a cluster. The file system can be a primary, secondary, or secondary-only. A stable image of the file system is provided for writes from any node.
- Multiple snapshots of a cluster file system can be mounted on the same or a different node in a cluster.
- A snapshot is accessible only on the node mounting a snapshot. The snapshot device cannot be mounted on two different nodes simultaneously.
- The device for mounting a snapshot can be a local disk or a shared volume. A shared volume is used exclusively by a snapshot mount and is not usable from other nodes in a cluster as long as the snapshot is active on that device.
- On the node mounting a snapshot, the snapped file system cannot be unmounted while the snapshot is mounted.
- A SFCFSHA snapshot ceases to exist if it is unmounted or the node mounting the snapshot fails. A snapshot, however, is not affected if any other node leaves or joins the cluster.
- A snapshot of a read-only mounted file system cannot be taken. It is possible to mount snapshot of a cluster file system only if the snapped cluster file system is mounted with the `crw` option.

Performance considerations

Mounting a snapshot file system for backup increases the load on the system because of the resources used to perform copy-on-writes and to read data blocks from the snapshot. In this situation, cluster snapshots can be used to do off-host backups. Off-host backups reduce the load of a backup application from the primary server. Overhead from remote snapshots is small when compared to overall snapshot overhead. Therefore, running a backup application by mounting a snapshot from a relatively less loaded node is beneficial to overall cluster performance.

Creating a snapshot on a Storage Foundation Cluster File System High Availability

To create and mount a snapshot on a two-node cluster using SFCFSHA administrative interface commands.

To create a snapshot on a cluster file system

- 1 To create a VxFS file system on a shared VxVM volume, type the following command:

```
# mkfs -F vxfs /dev/vx/rdisk/cfsdg/vol1
```

```
version 9 layout
```

```
104857600 sectors, 52428800 blocks of size 1024, log size 16384 blocks  
largefiles supported
```

- 2 To mount the file system on all nodes, type the following commands:

```
# cfsmntadm add cfsdg vol1 /mnt1 all=cluster
```

```
# cfsmount /mnt1
```

The `cfsmntadm` command adds an entry to the cluster manager configuration, then the `cfsmount` command mounts the file system on all nodes.

- 3 Add the snapshot on a previously created volume (`snapvol` in this example) to the cluster manager configuration. For example:

```
# cfsmntadm add snapshot cfsdg snapvol /mnt1 /mnt1snap system01=ro
```

The snapshot of a cluster file system is accessible only on the node where it is created; the snapshot file system itself cannot be cluster mounted.

- 4 Create and locally mount the snapshot file system on `system01`, type the following command:

```
# cfsmount /mnt1snap
```

- 5 A snapped file system cannot be unmounted until all of its snapshots are unmounted. Unmount and destroy the snapshot before trying to unmount the snapped cluster file system, type the following command:

```
# cfsunmount /mnt1snap
```

Administering VCS

This section provides instructions for the following VCS administration tasks:

- Viewing available Veritas devices and drivers
- Verifying VCS configuration
- Starting and stopping VCS

If you encounter issues while administering VCS, refer to the troubleshooting section for assistance.

Viewing available Veritas device drivers

To view the available Veritas devices:

```
# modinfo|grep -i vx
```

To view the devices that are loaded in memory, run the `modinfo` command as shown in the following examples.

For example:

If you want to view whether or not the driver 'gab' is loaded in memory:

```
# modinfo|grep gab
```

```
228 7b600000 53808 314 1 gab (GAB device 5.1SP1)
```

Configuring VCS to start Oracle with a specified Pfile

If you want to configure VCS such that Oracle starts with a specified Pfile, modify the `main.cf` file for the Oracle group as follows:

```
Oracle oral (
    Sid @system01 = vrts1
    Sid @system02 = vrts2
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = SRVCTLSTART
    ShutDownOpt = SRVCTLSTOP
    pfile="/app/oracle/orahome/dbs/initprod1.ora"
)
```

Verifying VCS configuration

To verify the VCS configuration:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

Starting and stopping VCS

To start VCS on each node:

```
# hstart
```

To stop VCS on each node:

```
# hstop -local
```

You can also use the command `hstop -all`; however, make sure that you wait for port 'h' to close before restarting VCS.

Configuring destination-based load balancing for LLT

Destination-based load balancing for LLT is turned off by default. Symantec recommends destination-based load balancing when the cluster setup has more than two nodes and more active LLT ports.

To configure destination-based load balancing for LLT

- ◆ Run the following command to configure destination-based load balancing:

```
lltconfig -F linkburst:0
```

Administering CVM

Listing all the CVM shared disks

You can use the following command to list all the CVM shared disks:

```
# vxdisk -o alldgs list |grep shared
```

Establishing CVM cluster membership manually

In most cases you do not have to start CVM manually; it normally starts when VCS is started.

Run the following command to start CVM manually:

```
# vxclustadm -m vcs -t gab startnode
```

```
vxclustadm: initialization completed
```

Note that `vxclustadm` reads `main.cf` for cluster configuration information and is therefore not dependent upon VCS to be running. You do not need to run the `vxclustadm startnode` command as normally the `hstart` (VCS start) command starts CVM automatically.

To verify whether CVM is started properly:

```
# vxclustadm nidmap
Name          CVM Nid    CM Nid     State
system01      0          0          Joined: Master
system02      1          1          Joined: Slave
```

Methods to control CVM master selection

When a master node leaves, Cluster Volume Manager (CVM) fails over the master role to another node in the cluster. CVM selects the node in the cluster that is best suited to take over the master role. CVM gives preference to nodes that have connectivity to the maximum number of disks in the disk group. This behavior is an enhancement over previous releases of CVM.

During regular operations, CVM dynamically assigns an offset preference value to each node. The preference assignment is automatic, and generally does not require any intervention from the storage administrator. However, if you need greater control over the master selection, you can also set customized preference values. When a master failover occurs, CVM uses the custom node preferences together with the offset preference values to select the new master node.

See [“About setting cluster node preferences for master failover”](#) on page 367.

To perform scheduled maintenance on a master node, you can manually migrate the master role to another node in the cluster.

See [“About changing the CVM master manually”](#) on page 372.

About setting cluster node preferences for master failover

Cluster Volume Manager (CVM) dynamically assigns an offset preference value to each node, depending on criteria such as the connectivity to the disks in the disk group. The preference assignment is automatic and generally does not require any intervention from the storage administrator.

If you need greater control over the master selection, you can set customized preference values. Determine which nodes in the CVM cluster are the most appropriate candidates to run the master role. Assign high or low preference values to the nodes so that CVM selects the master node from the pool of the most appropriate nodes. CVM uses the custom node preferences together with the offset preference values to select the new master node. CVM fails over the master role to the preferred node.

Cluster node preference for master failover

Cluster Volume Manager (CVM) assigns weight values to each node in the cluster based on internal criteria, such as the disk connectivity. CVM assigns weight values

from -100 to 100. A negative value means that the node is less likely to become the master node in case of failover. A positive value means that the node is more likely to become the master node.

If the CVM default values produce the desired behavior, you need not adjust the preference values. In some cases, you as administrator want more control over which nodes may become master during a failover. You can assign a positive value or a negative value as a custom weight on certain nodes. The customized preference values are static and persistent.

The custom weight values do not prevent CVM from assigning weights. CVM increases or decreases the weight values, starting from the custom value instead of starting from the default value of 0. To make sure that your preference has the effect that you want, consider the effect of the values that CVM generates.

For example, you want to ensure that NodeA always has preference for failover over NodeB. If NodeA loses connectivity to some disks, CVM reduces the weight value of NodeA. CVM also might increase the weight value of NodeB. To override the preference weights that CVM assigns, the custom preference value of NodeA must be at least 200 more than the custom value of NodeB. For example, CVM assigns a value of -100 to NodeA and a value of 100 to NodeB. If you want NodeB to be the master failover candidate in case NodeA loses connectivity to all disks, assign NodeA a value of 99.

Considerations for setting CVM node preferences

You can determine which nodes in the Cluster Volume Manager (CVM) cluster are the most appropriate candidates to run the master role. Assign high or low preference values to the nodes so that CVM selects the master node from the pool of the most appropriate nodes.

Set the preference value with the CVM resource agent or the `vxclustadm` command. The preference values are in the range from -2147483648 to 2147483647.

If you do not specify custom preferences, CVM gives preference to the node with the maximum visibility to the storage to become the CVM master node.

To set CVM preferences for master failover, the cluster must be cluster protocol version 110 or greater.

The following scenarios can indicate that you specify preference values:

- A cluster running an I/O intensive application
Set the preferences for master failover so that after a failover, the application runs on the master node. After the application crashes due to master node panic, the application fails over to the new master. During recovery, the failover does not incur the cost of exchanging messages between slave node and master

node if the I/O intensive application is co-located on the master. This behavior improves the failover and recovery time.

- A cluster where frequent administrative operations create high loads of internal I/Os.

Set the preferences for master failover so that after a failover, the application runs on a slave node.

Storage Foundation issues I/Os for administrative operations such as creating volumes or snapshots. In this release, VxVM throttles the administrative I/Os when the application I/O load is high. Throttling reduces the effect of Storage Foundation generated I/Os on the application I/Os.

If your environment requires frequent administrative operations, you can also set master failover preferences to minimize the effect of the administrative I/Os. Set lower preference values for master failover on the nodes to which the applications may failover.

- A cluster where nodes have different storage and processing capacities.

If your cluster contains some nodes with high processing capacity, you want CVM to prefer those nodes to serve as the master role and to run the applications. After considering application placement options, decide how to prioritize the cluster nodes. Assign high preference values to nodes that have higher processing capacity or better throughput.

On the other hand, your cluster may contain a few low-capacity nodes, which are used for regular backup of data or an internal low-priority job. These low-capacity nodes may not have visibility to all the shared storage so you do not want the master role to run on these nodes. Set negative preference values on these nodes so that CVM does not select them as CVM master over a more favorable candidate.

- A campus cluster where the master node should be on the same site as the application.

You may want to keep the master node co-located with the site where application is running. If you have defined a scheme where application failover happens within a site, you also want the CVM master to fail over to a node within the site. Set high preference values at all the nodes of a site, so that the master stays within the site during master failure. Currently, CVM does not automatically give preference based on the site.

Setting the cluster node preference using the CVMCluster agent

On a cluster node, you can set the preference values for master failover using the CVMCluster agent. Preferences set with this method are persistent across reboots.

To set the cluster node preference using the CVMCluster agent

- 1 Make the configuration writable.

```
# haconf -makerw
```

- 2 View the preferences of the node at the current time:

```
# hares -display cvm_clus -attribute CVMNodePreference
```

The command displays an integer value specifying the preference of the local node to become the next master. Run the command on all the nodes and then check to see which node has the highest preference to become the next master.

- 3 Set new preference values:

```
# hares -modify cvm_clus CVMNodePreference \  
"node1=weight1, node2=weight2, ..."
```

Setting the cluster node preference value for master failover using the vxclustadm command

The procedures in this section describe how to set the cluster node preference value for master failover using the `vxclustadm` command. Preferences set with this method are not persistent across reboots.

To set the cluster node preference using the vxclustadm command

- 1 To view the existing preferences, use the following command:

```
# vxclustadm getpreference
```

- 2 To set a new preference, use the following command:

```
# vxclustadm setpreference value
```

Example of setting the cluster node preference value for master failover

This example describes a case where the cluster environment includes three types of nodes. You can define master preferences for each type of node, so that CVM fails over the master role to the best candidate.

Consider the three types of nodes as the following pools:

- Pool1 (Node1, Node2, Node3)


These nodes have high capacity (storage and processing) and have full storage visibility. An I/O intensive application is running at these nodes of the cluster. You are OK to choose a node from this pool over nodes from the other pool even if it loses 30% of its disks.

- Pool2 (Node4, Node5)**
 The cluster has a few low capacity (storage and processing) nodes. Internal (in-house) applications use these nodes to do post-processing over the data. Some of these nodes are also used for snapshots and backups. You want to choose a node from this pool:
 If there are no nodes to choose from Pool1 OR
 If all the nodes in Pool1 have lost access to substantial disks reducing their preference values below the preference values of nodes from Pool2.
- Pool3 (Node6, Node7)**
 These nodes run applications that do not need to work on all the volumes. These nodes may have a restricted view to the storage (limited storage visibility). CVM internally offsets the preference values of nodes in this pool, because they do not see all of the storage. To reduce the likelihood that one of these nodes becomes the master node, you can assign negative preference values.

If you do not define any custom preferences, CVM determines the preference as an offset depending on disk availability. Suppose at time A, the current master, Node1, leaves the cluster. CVM calculates the preferences, and chooses Node5 since it has the best connectivity. If Node5 leaves, CVM selects Node4. Nodes from Pool2 are preferred over nodes in Pool1.

Pool	Node	CVM offset at time A
Pool1	Node1 (Current master)	0
Pool1	Node2	-30
Pool1	Node3	-25
Pool2	Node4.	-20
Pool2	Node5	0
Pool3	Node6	-50
Pool3	Node7	-50

In this example, you want to specify higher preference values to nodes in Pool1 and lower preferences to the nodes in Pool3. The following diagram show possible preference values to the nodes in the cluster:



To set the cluster node preference

- 1 For each node in Pool1, set the preference to 30.

```
# hares -modify cvm_clus CVMNodePreference \  
"node1=30, node2=30, node3=30"
```

- 2 For each node in Pool3, set the preference to -50.

```
# hares -modify cvm_clus CVMNodePreference "node6=-50, node7=-50"
```

After you set the preference values as above, the behavior reflects the desired behavior in case of failover. If Node1 fails, the other nodes in Pool1 are the most likely candidates to take over the master. Node3 has lost 25 percent of its disks, but it is still preferred over the nodes in the other pools. Although Node5 has good connectivity, nodes in the Pool1 are preferred over this one.

Node	CVM offset at time A	Custom preference	Total preference
Node1 (current master)	0	30	30
Node2	-29	30	1
Node3	-25	30	5
Node4.	-20	0	-20
Node5	0	0	0
Node6	-50	-50	-100
.Node7	-50	-50	-100

About changing the CVM master manually

When you migrate the master role manually, you must specify the node that you want to take over the master role. You can view the preference values set on the node to help you determine the best node to take over the master role.

When you change the master node manually, the cluster stays online, and none of the nodes are required to leave the cluster. However, CVM quiesces the application I/Os. Therefore, schedule the master switch operation at an appropriate time to minimize the effect on performance.

After a master failover, if you decide that CVM selected an inappropriate node to be the master, you can change the master role manually. In this case, you may also want to change the customized failover preference values for the cluster nodes.

See [“About setting cluster node preferences for master failover”](#) on page 367.

Considerations for changing the master manually

If the master is not running on the node best suited to be the master of the cluster, you can manually change the master. Here are some scenarios when this might occur.

- The currently running master lost access to some of its disks.
By default, CVM uses I/O shipping to handle this scenario. However, you may want to failover the application to a node which has access to the disks. When you move the application, you may also want to relocate the master role to a new node. For example, you may want the master node and the application to be on the same node.
You can use the master switching operation to move the master role without causing the original master node to leave the cluster. After the master role and the application are both switched to other nodes, you may want to remove the original node from the cluster. You can unmount the file systems and cleanly shut down the node. You can then do maintenance on the node.
- The master node is not scaling well with the overlap of application load and the internally-generated administrative I/Os.
You may choose to reevaluate the placement strategy and relocate the master node.

See [“Considerations for setting CVM node preferences”](#) on page 368.

Changing the CVM master manually

You can change the CVM master manually from one node in the cluster to another node, while the cluster is online. CVM migrates the master node, and reconfigures the cluster.

Symantec recommends that you switch the master when the cluster is not handling VxVM configuration changes or cluster reconfiguration operations. In most cases, CVM aborts the operation to change the master, if CVM detects that any configuration changes are occurring in the VxVM or the cluster. After the master

change operation starts reconfiguring the cluster, other commands that require configuration changes will fail until the master switch completes.

See “Errors during CVM master switching” on page 375.

To change the master online, the cluster must be cluster protocol version 100 or greater.

To change the CVM master manually

- 1 To view the current master, use one of the following commands:

```
# vxclustadm nidmap
Name          CVM Nid    CM Nid    State
system01      0          0         Joined: Slave
system02      1          1         Joined: Master

# vxdctl -c mode
mode: enabled: cluster active - MASTER
master: system02
```

In this example, the CVM master is system02.

- 2 From any node on the cluster, run the following command to change the CVM master:

```
# vxclustadm setmaster nodename
```

where *nodename* specifies the name of the new CVM master.

The following example shows changing the master on a cluster from system02 to system01:

```
# vxclustadm setmaster system01
```

- 3 To monitor the master switching, use the following command:

```
# vxclustadm -v nodestate
state: cluster member
      nodeId=0
      masterId=0
      neighborId=1
      members[0]=0xf
      joiners[0]=0x0
      leavers[0]=0x0
      members[1]=0x0
      joiners[1]=0x0
      leavers[1]=0x0
      reconfig_seqnum=0x9f9767
      vxfen=off
state: master switching in progress
reconfig: vxconfigd in join
```

In this example, the state indicates that master is being changed.

- 4 To verify whether the master has successfully changed, use one of the following commands:

```
# vxclustadm nidmap
Name          CVM Nid    CM Nid    State
system01      0          0         Joined: Master
system02      1          1         Joined: Slave

# vxdctl -c mode
mode: enabled: cluster active - MASTER
master: system01
```

Errors during CVM master switching

Symantec recommends that you switch the master when the cluster is not handling VxVM or cluster configuration changes.

In most cases, CVM aborts the operation to change the master, if CVM detects any configuration changes in progress. CVM logs the reason for the failure into the system logs. In some cases, the failure is displayed in the `vxclustadm setmaster` output as follows:

```
# vxclustadm setmaster system01
VxVM vxclustadm ERROR V-5-1-0 Master switching, a reconfiguration or
a transaction is in progress.
Try again
```

In some cases, if the master switching operation is interrupted with another reconfiguration operation, the master change fails. In this case, the existing master remains the master of the cluster. After the reconfiguration is complete, reissue the `vxclustadm setmaster` command to change the master.

If the master switching operation has started the reconfiguration, any command that initiates a configuration change fails with the following error:

```
Node processing a master-switch request. Retry operation.
```

If you see this message, retry the command after the master switching has completed.

Importing a shared disk group manually

You can use the following command to manually import a shared disk group:

```
# vxdbg -s import dg_name
```

Deporting a shared disk group manually

You can use the following command to manually deport a shared disk group:

```
# vxdbg deport dg_name
```

Note that the deport of a shared disk group removes the SCSI-3 PGR keys on the disks.

Starting shared volumes manually

Following a manual CVM shared disk group import, the volumes in the disk group need to be started manually, as follows:

```
# vxvol -g dg_name startall
```

To verify that the volumes are started, run the following command:

```
# vxprint -htrg dg_name | grep ^v
```

Evaluating the state of CVM ports

CVM kernel (vxio driver) uses port 'v' for kernel messaging, port 'w' for vxconfigd communication between the cluster nodes, and port 'u' for shipping commands from the slave node to the master node. The following command displays the state of CVM ports:

```
# gabconfig -a | egrep "Port [vwu]"
```

Verifying if CVM is running in an SFCFSHA cluster

You can use the following options to verify whether CVM is up or not in an SFCFSHA cluster.

The following output is displayed on a node that is not a member of the cluster:

```
# vxctl -c mode
mode: enabled: cluster inactive
# vxclustadm -v nodestate
state: out of cluster
```

On the master node, the following output is displayed:

```
# vxctl -c mode
mode: enabled: cluster active - MASTER
master: system01
```

On the slave nodes, the following output is displayed:

```
# vxctl -c mode
mode: enabled: cluster active - SLAVE
master: system02
```

The following command lets you view all the CVM nodes at the same time:

```
# vxclustadm nidmap
```

Name	CVM Nid	CM Nid	State
system01	0	0	Joined: Master
system02	1	1	Joined: Slave

Verifying CVM membership state

The state of CVM can be verified as follows:

```
# vxclustadm -v nodestate
state: joining
      nodeId=0
      masterId=0
      neighborId=1
      members=0x3
      joiners=0x0
```

```
leavers=0x0
reconfig_segnum=0x72a10b
vxfen=on
```

The state indicates that CVM has completed its kernel level join and is in the middle of vxconfigd level join.

The `vxdtctl -c mode` command indicates whether a node is a CVM master or CVM slave.

Verifying the state of CVM shared disk groups

You can use the following command to list the shared disk groups currently imported in the SFCFSHA cluster:

```
# vxdg list |grep shared

oradatadg enabled,shared 1052685125.1485.system01
```

Verifying the activation mode

In an SFCFSHA cluster, the activation of shared disk group should be set to “shared-write” on each of the cluster nodes.

To verify whether the “shared-write” activation is set:

```
# vxdg list dg_name |grep activation

local-activation: shared-write
```

If “shared-write” activation is not set, run the following command:

```
# vxdg -g dg_name set activation=sw
```

CVM log files

The `/var/VRTSvcs/log` directory contains the agent log files.

```
# cd /var/VRTSvcs/log
# ls -l *CVM* engine_A.log
CVMCluster_A.log      # CVM Agent log
CVMVolDg_A.log        # CVM VolDg Agent log
CVMVxconfigd_A.log    # CVM vxconfigd Agent log
engine_A.log          # VCS log
```

You can use the `cmdlog` file to view the list of CVM commands that have been executed. The file is located at `/var/adm/vx/cmdlog`.

Requesting node status and discovering the master node

The `vxctl` utility controls the operation of the `vxconfigd` volume configuration daemon. The `-c` option can be used to request cluster information and to find out which node is the master. To determine whether the `vxconfigd` daemon is enabled and/or running, use the following command:

```
vxctl -c mode
```

[Table 15-2](#) shows the various messages that may be output according to the current status of the cluster node.

Table 15-2 Cluster status messages

Status message	Description
mode: enabled: cluster active - MASTER master: mozart	The node is the master.
mode: enabled: cluster active - SLAVE master: mozart	The node is a slave.
mode: enabled: cluster active - role not set master: mozart state: joining reconfig: master update	The node has not yet been assigned a role, and is in the process of joining the cluster.
mode: enabled: cluster active - SLAVE master: mozart state: joining	The node is configured as a slave, and is in the process of joining the cluster.
mode: enabled: cluster inactive	The cluster is not active on this node.
mode: booted: master: ts4200-04	Enable root disk encapsulation but not transactions.
mode: disabled:	Disable transactions.

If the `vxconfigd` daemon is disabled, no cluster information is displayed.

See the `vxdctl(1M)` manual page.

Determining if a LUN is in a shareable disk group

The `vxdisk` utility manages VxVM disks. To use the `vxdisk` utility to determine whether a LUN is part of a cluster-shareable disk group, use the following command:

```
# vxdisk list accessname
```

where *accessname* is the disk access name (or device name).

For example, a portion of the output from this command (for the device `c4t1d0`) is shown here:

```
Device:      c4t1d0
devicetag:   c4t1d0
type:        auto
clusterid:   cvm2
disk:        name=shdg01 id=963616090.1034.cvm2
timeout:     30
group:       name=shdg id=963616065.1032.cvm2
flags:       online ready autoconfig shared imported
...
```

Note that the `clusterid` field is set to `cvm2` (the name of the cluster), and the `flags` field includes an entry for `shared`. The `imported` flag is only set if a node is a part of the cluster and the disk group is imported.

Listing shared disk groups

`vxdg` can be used to list information about shared disk groups. To display information for all disk groups, use the following command:

```
# vxdg list
```

Example output from this command is displayed here:

NAME	STATE	ID
group2	enabled,shared	774575420.1170.teal
group1	enabled,shared	774222028.1090.teal

Shared disk groups are designated with the flag `shared`.

To display information for shared disk groups only, use the following command:

```
# vxdg -s list
```


Example output from this command is as follows:

NAME	STATE	ID
group2	enabled,shared	774575420.1170.teal
group1	enabled,shared	774222028.1090.teal

To display information about one specific disk group, use the following command:

```
# vxdg list diskgroup
```

The following is example output for the command `vxdg list tempdg` on the master:

```
Group:      tempdg
dgid:       1245902808.74.ts4200-04
import-id:  33792.73
flags:      shared cds
version:    150
alignment:  8192 (bytes)
local-activation: shared-write
cluster-actv-modes: ts4200-04=sw ts4200-06=sw ts4200-05=sw
ssb:        on
autotagging: on
detach-policy: global
dg-fail-policy: dgdisable
copies:     nconfig=default nlog=default
config:     seqno=0.1027 permten=0 free=0 templen=0 loglen=0
```

Note that the `flags` field is set to `shared`. The output for the same command when run on a slave is slightly different. The `local-activation` and `cluster-actv-modes` fields display the activation mode for this node and for each node in the cluster respectively. The `detach-policy` and `dg-fail-policy` fields indicate how the cluster behaves in the event of loss of connectivity to the disks, and to the configuration and log copies on the disks.

Creating a shared disk group

You can run the command to create a shared disk group on a master node or a slave node. If you create the disk group on a slave node, the command is shipped to the master and executed on the master.

If the cluster software has been run to set up the cluster, a shared disk group can be created using the following command:

```
# vxdg -s init diskgroup [diskname=]devicenames
```

where *diskgroup* is the disk group name, *diskname* is the administrative name chosen for a VM disk, and *devicename* is the device name (or disk access name).

Warning: The operating system cannot tell if a disk is shared. To protect data integrity when dealing with disks that can be accessed by multiple systems, use the correct designation when adding a disk to a disk group. VxVM allows you to add a disk that is not physically shared to a shared disk group if the node where the disk is accessible is the only node in the cluster. However, this means that other nodes cannot join the cluster. Furthermore, if you attempt to add the same disk to different disk groups (private or shared) on two nodes at the same time, the results are undefined. Perform all configuration on one node only, and preferably on the master node.

Importing disk groups as shared

You can import shared disk groups on a master node or a slave node. If you run the command to import the shared disk group on a slave node, the command is shipped to the master and executed on the master.

Disk groups can be imported as shared using the `vxvg -s import diskgroup` command. If the disk groups are set up before the cluster software is run, the disk groups can be imported into the cluster arrangement using the following command:

```
# vxvg -s import diskgroup
```

where *diskgroup* is the disk group name or ID. On subsequent cluster restarts, the disk group is automatically imported as shared. Note that it can be necessary to deport the disk group (using the `vxvg deport diskgroup` command) before invoking the `vxvg` utility.

Forcibly importing a disk group

You can use the `-f` option to the `vxvg` command to import a disk group forcibly.

Warning: The force option (`-f`) must be used with caution and only if you are fully aware of the consequences such as possible data corruption.

When a cluster is restarted, VxVM can refuse to auto-import a disk group for one of the following reasons:

- A disk in the disk group is no longer accessible because of hardware errors on the disk. In this case, use the following command to forcibly reimport the disk group:

```
# vxdg -s -f import diskgroup
```

Note: After a forced import, the data on the volumes may not be available and some of the volumes may be in the disabled state.

- Some of the disks in the shared disk group are not accessible, so the disk group cannot access all of its disks. In this case, a forced import is unsafe and must not be attempted because it can result in inconsistent mirrors.

Converting a disk group from shared to private

You can convert shared disk groups on a master node or a slave node. If you run the command to convert the shared disk group on a slave node, the command is shipped to the master and executed on the master.

To convert a shared disk group to a private disk group, first deport it on the master node using this command:

```
# vxdg deport diskgroup
```

Then reimport the disk group on any cluster node using this command:

```
# vxdg import diskgroup
```

Moving objects between shared disk groups

You can move objects between shared disk groups on a master node or a slave node. If you run the command to move objects between shared disk groups on a slave node, the command is shipped to the master and executed on the master.

You can use the `vxdg move` command to move a self-contained set of VxVM objects such as disks and top-level volumes between disk groups. In a cluster, you can move such objects between private disk groups on any cluster node where those disk groups are imported.

See [“Moving objects between disk groups”](#) on page 821.

Splitting shared disk groups

You can use the `vxdg split` command to remove a self-contained set of VxVM objects from an imported disk group, and move them to a newly created disk group.

See [“Splitting disk groups”](#) on page 824.

Splitting a private disk group creates a private disk group, and splitting a shared disk group creates a shared disk group. You can split a private disk group on any cluster node where that disk group is imported.

You can split a shared disk group or create a shared target disk group on a master node or a slave node. If you run the command to split a shared disk group or to create a shared target disk group on a slave node, the command is shipped to the master and executed on the master.

Joining shared disk groups

You cannot join a private disk group and a shared disk group.

You can use the `vxdg join` command to merge the contents of two imported disk groups. In a cluster, you can join two private disk groups on any cluster node where those disk groups are imported.

If the source disk group and the target disk group are both shared, you can perform the join from a master node or a slave node. If you run the command to perform the join on a slave node, the command is shipped to the master and executed on the master.

See [“Joining disk groups”](#) on page 825.

Changing the activation mode on a shared disk group

The activation mode for access by a cluster node to a shared disk group is set directly on that node.

The activation mode of a shared disk group can be changed using the following command:

```
# vxdg -g diskgroup set activation=mode
```

The activation mode is one of `exclusivewrite` or `ew`, `readonly` or `ro`, `sharedread` or `sr`, `sharedwrite` or `sw`, or `off`.

If you use this command to change the activation mode of a shared disk group, you must first change the activation mode to `off` before setting it to any other value, as shown here:

```
# vxdg -g myshdg set activation=off
# vxdg -g myshdg set activation=readonly
```

See [“Activation modes of shared disk groups”](#) on page 150.

Setting the detach policy for shared disk groups

The default for the detach policy is `global`.

To change the detach policy on a shared disk group

- ◆ To set the detach policy to local on a shared disk group.

```
# vxdbg -g diskgroup set diskdetpolicy=local
```

To set the detach policy to global (the default).

```
# vxdbg -g diskgroup set diskdetpolicy=global
```

Handling cloned disks in a shared disk group

If a disk is cloned or copied in such a way to create a duplicate disk ID, you must perform special actions to import the disk into VxVM. The procedures are the same for shared disk groups as for private disk groups. When you are ready to import the disk, specify the `-s` option to the `vxdbg import` command:

```
# vxdbg import -s diskgroup
```

Creating volumes with exclusive open access by a node

When using the `vxassist` command to create a volume, you can use the `exclusive=on` attribute to specify that the volume may only be opened by one node in the cluster at a time. For example, to create the mirrored volume `volmir` in the disk group `dskgrp`, and configure it for exclusive open, use the following command:

```
# vxassist -g dskgrp make volmir 5g layout=mirror exclusive=on
```

Multiple opens by the same node are also supported. Any attempts by other nodes to open the volume fail until the final close of the volume by the node that opened it.

Specifying `exclusive=off` instead means that more than one node in a cluster can open a volume simultaneously. This is the default behavior.

Setting exclusive open access to a volume by a node

Exclusive open access on a volume can be set from the any node in the cluster. Ensure that none of the nodes in the cluster have the volume open when setting this attribute.

You can set the `exclusive=on` attribute with the `vxvol` command to specify that an existing volume may only be opened by one node in the cluster at a time.

For example, to set exclusive open on the volume `volmir` in the disk group `dskgrp`, use the following command:

```
# vxvol -g dskgrp set exclusive=on volmir
```

Multiple opens by the same node are also supported. Any attempts by other nodes to open the volume fail until the final close of the volume by the node that opened it.

Specifying `exclusive=off` instead means that more than one node in a cluster can open a volume simultaneously. This is the default behavior.

Displaying the cluster protocol version

The following command displays the cluster protocol version running on a node:

```
# vxdctl list
```

This command produces output similar to the following:

```
Volboot file
version: 3/1
seqno: 0.19
cluster protocol version: 110
hostid: giga
entries:
```

You can also check the existing cluster protocol version using the following command:

```
# vxdctl protocolversion
```

This produces output similar to the following:

```
Cluster running at protocol 110
```

Displaying the supported cluster protocol version range

The following command displays the maximum and minimum protocol version supported by the node and the current protocol version:

```
# vxdctl support
```

This command produces output similar to the following:

```
Support information:
vxconfigd_vrsn:      31
dg_minimum:          20
dg_maximum:          170
kernel:              31
```

```
protocol_minimum: 100
protocol_maximum: 110
protocol_current: 110
```

You can also use the following command to display the maximum and minimum cluster protocol version supported by the current Veritas Volume Manager release:

```
# vxctl protocolrange
```

This produces output similar to the following:

```
minprotoversion: 100, maxprotoversion: 110
```

Recovering volumes in shared disk groups

The `vxrecover` utility is used to recover plexes and volumes after disk replacement. When a node leaves a cluster, it can leave some mirrors in an inconsistent state. The `vxrecover` utility can be used to recover such volumes. The `-c` option to `vxrecover` causes it to recover all volumes in shared disk groups. The `vxconfigd` daemon automatically calls the `vxrecover` utility with the `-c` option when necessary.

Warning: While the `vxrecover` utility is active, there can be some degradation in system performance.

Obtaining cluster performance statistics

The `vxstat` utility returns statistics for specified objects. In a cluster environment, `vxstat` gathers statistics from all of the nodes in the cluster. The statistics give the total usage, by all nodes, for the requested objects. If a local object is specified, its local usage is returned.

You can optionally specify a subset of nodes using the following form of the command:

```
# vxstat -g diskgroup -n node[,node...]
```

where `node` is the CVM node ID number. You can find out the CVM node ID by using the following command:

```
# vxclustadm nidmap
```

If a comma-separated list of nodes is supplied, the `vxstat` utility displays the sum of the statistics for the nodes in the list.

For example, to obtain statistics for node 2, volume `vol1`, use the following command:

```
# vxstat -g diskgroup -n 2 vol1
```

This command produces output similar to the following:

		OPERATIONS		BLOCKS		AVG TIME (ms)	
TYP	NAME	READ	WRITE	READ	WRITE	READ	WRITE
vol	vol1	2421	0	600000	0	99.0	0.0

To obtain and display statistics for the entire cluster, use the following command:

```
# vxstat -b
```

The statistics for all nodes are summed. For example, if node 1 performed 100 I/O operations and node 2 performed 200 I/O operations, `vxstat -b` displays a total of 300 I/O operations.

Administering CVM from the slave node

CVM requires that the master node of the cluster executes configuration commands, which change the object configuration of a CVM shared disk group. Examples of configuration changes include creating shared disk groups, importing shared disk groups, deporting shared disk groups, and creating volumes or snapshots in a shared disk group.

Starting in this release, you can issue most configuration commands that operate on the shared disk group from any node in the cluster. If you issue the command on the slave node, CVM ships the commands from the slave node to the master node. CVM then executes the command on the master node. In normal conditions, we recommend that you issue configuration-changing commands for a shared disk group on the master node. If the circumstances require, you can issue these commands from the slave node.

Commands that operate on private disk groups are not shipped to the master node. Similarly, CVM does not ship commands that operate locally on the slave node, such as `vxprint` and `vxdisk list`.

CVM uses the Group Membership Services and Atomic Broadcast (GAB) transport mechanism of Veritas Cluster Server (VCS) to ship the commands from the slave node to the master node.

When you issue a command on the slave that is executed on the master, the command output (on the slave node) displays the object names corresponding to the master node. For example, the command displays the disk access name (`daname`) from the master node.

When run from a slave node, a query command such as `vxtask` or `vxstat` displays the status of the commands on the slave node. The command does not show the

status of commands that originated from the slave node and that are executing on the master node.

Note the following error handling for commands that you originate from the slave node, which CVM executes on the master:

- If the `vxconfigd` daemon on either the slave node or on the master node fails, the command exits. The instance of the command on the master also exits. To determine if the command executed successfully, use the `vxprint` command to check the status of the VxVM objects.
- If the slave node that shipped the command or the master node leaves the cluster while the master is executing the command, the command exits on the master node as well as on the slave node. To determine if the command executed successfully, use the `vxprint` command to check the status of the VxVM objects.

Note the following limitations for issuing CVM commands from the slave node:

- The CVM protocol version must be at least 100 on all nodes in the cluster. See [“Displaying the cluster protocol version”](#) on page 386.
- CVM uses the values in the defaults file on the master node when CVM executes the command. To avoid any ambiguity, we recommend that you use the same values in the defaults file for each of the nodes in the cluster.
- CVM does not support executing all commands on the slave node. You must issue the following commands only on the master node:

- Commands that specify a controller name. For example:

```
# vxassist -g shareddg make sharedvol 20M ctrl:fscsi0
```

- Commands that specify both a shared disk group and a private disk group. For example:

```
# vxdg destroy privatedg shareddg
```

- Commands that include the defaults file as an argument. For example:

```
# vxassist -d defaults_file
```

- Veritas Volume Replicator (VVR) commands including `vxibc`, `vxrlink`, `vxrsync`, `vxrvlg`, `vrport`, `vrstat`, and `vradmin`.

- The `vxdisk` command options that act on shared disk groups.

See [“CVM commands supported for executing on the slave node”](#) on page 964.

Administering ODM

This section provides instructions for the following ODM administration tasks:

- [Verifying the ODM port](#)
- [Starting ODM](#)

If you encounter issues while administering ODM, refer to the troubleshooting section for assistance.

Verifying the ODM port

It is recommended to enable ODM in SFCFSHA. Run the following command to verify that ODM is running:

```
# gabconfig -a | grep "Port d"
```

Starting ODM

The following procedure provides instructions for starting ODM.

To start ODM

- ◆ Run the following command:

For Solaris 10:

```
# svcadm enable vxodm
```

Administering I/O Fencing

See the *Veritas Cluster Server Administrator's Guide* for more information.

About administering I/O fencing

The I/O fencing feature provides the following utilities that are available through the VRTSvxfen package:

vxfcntlshdw	Tests SCSI-3 functionality of the disks for I/O fencing See “About the vxfcntlshdw utility” on page 391.
vxfenconfig	Configures and unconfigures I/O fencing Lists the coordination points used by the vxfen driver.

vxfenadm	Displays information on I/O fencing operations and manages SCSI-3 disk registrations and reservations for I/O fencing See “About the vxfenadm utility” on page 399.
vxfenclearpre	Removes SCSI-3 registrations and reservations from disks See “About the vxfenclearpre utility” on page 404.
vxfenswap	Replaces coordination points without stopping I/O fencing See “About the vxfenswap utility” on page 407.
vxfendisk	Generates the list of paths of disks in the diskgroup. This utility requires that Veritas Volume Manager is installed and configured.

The I/O fencing commands reside in the /opt/VRTS/bin folder. Make sure you added this folder path to the PATH environment variable.

Refer to the corresponding manual page for more information on the commands.

About the vxfentsthdw utility

You can use the vxfentsthdw utility to verify that shared storage arrays to be used for data support SCSI-3 persistent reservations and I/O fencing. During the I/O fencing configuration, the testing utility is used to test a single disk. The utility has other options that may be more suitable for testing storage devices in other configurations. You also need to test coordinator disk groups.

See *Veritas Storage Foundation Cluster File System High Availability Installation Guide* to set up I/O fencing.

The utility, which you can run from one system in the cluster, tests the storage used for data by setting and verifying SCSI-3 registrations on the disk or disks you specify, setting and verifying persistent reservations on the disks, writing data to the disks and reading it, and removing the registrations from the disks.

Refer also to the vxfentsthdw(1M) manual page.

About general guidelines for using vxfentsthdw utility

Review the following guidelines to use the vxfentsthdw utility:

- The utility requires two systems connected to the shared storage.

Caution: The tests overwrite and destroy data on the disks, unless you use the `-r` option.

- The two nodes must have ssh (default) or rsh communication. If you use rsh, launch the `vxfcntlsthdw` utility with the `-n` option.
After completing the testing process, you can remove permissions for communication and restore public network connections.
- To ensure both systems are connected to the same disk during the testing, you can use the `vxfenadm -i diskpath` command to verify a disk's serial number. See [“Verifying that the nodes see the same disk”](#) on page 403.
- For disk arrays with many disks, use the `-m` option to sample a few disks before creating a disk group and using the `-g` option to test them all.
- The utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rdisk/clt1d0s2 is ready to be configured for
I/O Fencing on node system01
```

If the utility does not show a message stating a disk is ready, verification has failed.

- If the disk you intend to test has existing SCSI-3 registration keys, the test issues a warning before proceeding.

About the vxfcntlsthdw command options

[Table 15-3](#) describes the methods that the utility provides to test storage devices.

Table 15-3 vxfcntlsthdw options

vxfcntlsthdw option	Description	When to use
-n	Utility uses rsh for communication.	Use when rsh is used for communication.
-r	Non-destructive testing. Testing of the disks for SCSI-3 persistent reservations occurs in a non-destructive way; that is, there is only testing for reads, not writes. May be used with -m, -f, or -g options.	Use during non-destructive testing. See “Performing non-destructive testing on the disks using the -r option” on page 395.
-t	Testing of the return value of SCSI TEST UNIT (TUR) command under SCSI-3 reservations. A warning is printed on failure of TUR testing.	When you want to perform TUR testing.

Table 15-3 vxfcntlsthdw options (*continued*)

vxfcntlsthdw option	Description	When to use
-d	Use DMP devices. May be used with -c or -g options.	By default, the script picks up the DMP paths for disks in the disk group. If you want the script to use the raw paths for disks in the disk group, use the -w option.
-w	Use raw devices. May be used with -c or -g options.	With the -w option, the script picks the operating system paths for disks in the disk group. By default, the script uses the -d option to pick up the DMP paths for disks in the disk group.
-c	Utility tests the coordinator disk group prompting for systems and devices, and reporting success or failure.	For testing disks in coordinator disk group. See “Testing the coordinator disk group using vxfcntlsthdw -c option” on page 394.
-m	Utility runs manually, in interactive mode, prompting for systems and devices, and reporting success or failure. May be used with -r and -t options. -m is the default option.	For testing a few disks or for sampling disks in larger arrays. See “Testing the shared disks using the vxfcntlsthdw -m option” on page 395.
-f <i>filename</i>	Utility tests system/device combinations listed in a text file. May be used with -r and -t options.	For testing several disks. See “Testing the shared disks listed in a file using the vxfcntlsthdw -f option” on page 397.
-g <i>disk_group</i>	Utility tests all disk devices in a specified disk group. May be used with -r and -t options.	For testing many disks and arrays of disks. Disk groups may be temporarily created for testing purposes and destroyed (ungrouped) after testing. See “Testing all the disks in a disk group using the vxfcntlsthdw -g option” on page 398.

Testing the coordinator disk group using vxfcntlsthaw -c option

Use the vxfcntlsthaw utility to verify disks are configured to support I/O fencing. In this procedure, the vxfcntlsthaw utility tests the three disks one disk at a time from each node.

The procedure in this section uses the following disks for example:

- From the node system01, the disks are seen as /dev/rdisk/c1t1d0s2, /dev/rdisk/c2t1d0s2, and /dev/rdisk/c3t1d0s2.
- From the node system02, the same disks are seen as /dev/rdisk/c4t1d0s2, /dev/rdisk/c5t1d0s2, and /dev/rdisk/c6t1d0s2.

Note: To test the coordinator disk group using the vxfcntlsthaw utility, the utility requires that the coordinator disk group, vxfcntlcoordg, be accessible from two nodes.

To test the coordinator disk group using vxfcntlsthaw -c

- 1 Use the vxfcntlsthaw command with the -c option. For example:

```
# vxfcntlsthaw -c vxfcntlcoordg
```

- 2 Enter the nodes you are using to test the coordinator disks:

```
Enter the first node of the cluster: system01
```

```
Enter the second node of the cluster: system02
```

- 3 Review the output of the testing process for both nodes for all disks in the coordinator disk group. Each disk should display output that resembles:

```
ALL tests on the disk /dev/rdisk/c1t1d0s2 have PASSED.  
The disk is now ready to be configured for I/O Fencing on node  
system01 as a COORDINATOR DISK.
```

```
ALL tests on the disk /dev/rdisk/c4t1d0s2 have PASSED.  
The disk is now ready to be configured for I/O Fencing on node  
system02 as a COORDINATOR DISK.
```

- 4 After you test all disks in the disk group, the vxfcntlcoordg disk group is ready for use.

Removing and replacing a failed disk

If a disk in the coordinator disk group fails verification, remove the failed disk or LUN from the vxencoordg disk group, replace it with another, and retest the disk group.

To remove and replace a failed disk

- 1 Use the vxdiskadm utility to remove the failed disk from the disk group.
Refer to the *Veritas Storage Foundation Administrator's Guide*.
- 2 Add a new disk to the node, initialize it, and add it to the coordinator disk group.
See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for instructions to initialize disks for I/O fencing and to set up coordinator disk groups.
If necessary, start the disk group.
See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for instructions to start the disk group.
- 3 Retest the disk group.
See [“Testing the coordinator disk group using vxfcntlshdw -c option”](#) on page 394.

Performing non-destructive testing on the disks using the -r option

You can perform non-destructive testing on the disk devices when you want to preserve the data.

To perform non-destructive testing on disks

- ◆ To test disk devices containing data you want to preserve, you can use the -r option with the -m, -f, or -g options.
For example, to use the -m option and the -r option, you can run the utility as follows:

```
# vxfcntlshdw -rm
```

When invoked with the -r option, the utility does not use tests that write to the disks. Therefore, it does not test the disks for all of the usual conditions of use.

Testing the shared disks using the vxfcntlshdw -m option

Review the procedure to test the shared disks. By default, the utility uses the -m option.

This procedure uses the `/dev/rdisk/c1t1d0s2` disk in the steps.

If the utility does not show a message stating a disk is ready, verification has failed. Failure of verification can be the result of an improperly configured disk array. It can also be caused by a bad disk.

If the failure is due to a bad disk, remove and replace it. The `vxfcntlsthaw` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rdisk/c1t1d0s2 is ready to be configured for
I/O Fencing on node system01
```

Note: For A/P arrays, run the `vxfcntlsthaw` command only on active enabled paths.

To test disks using `vxfcntlsthaw` script

- 1 Make sure system-to-system communication is functioning properly.
- 2 From one node, start the utility.

```
# vxfcntlsthaw [-n]
```

- 3 After reviewing the overview and warning that the tests overwrite data on the disks, confirm to continue the process and enter the node names.

```
***** WARNING!!!!!!!!!! *****
```

```
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!
```

```
Do you still want to continue : [y/n] (default: n) y
```

```
Enter the first node of the cluster: system01
```

```
Enter the second node of the cluster: system02
```


- 4 Enter the names of the disks you are checking. For each node, the disk may be known by the same name:

```
Enter the disk name to be checked for SCSI-3 PGR on node
system01 in the format:
```

```
for dmp: /dev/vx/rdmp/cxtxdxs2
```

```
for raw: /dev/rdisk/cxtxdxs2
```

```
/dev/rdsk/c2t13d0s2
```

```
Make sure it's the same disk as seen by nodes system01 and system02
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
system02 in the format:
```

```
for dmp: /dev/vx/rdmp/cxtxdxs2
```

```
for raw: /dev/rdisk/cxtxdxs2
```

```
Make sure it's the same disk as seen by nodes system01 and system02
```

```
/dev/rdsk/c2t13d0s2
```

If the serial numbers of the disks are not identical, then the test terminates.

- 5 Review the output as the utility performs the checks and report its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success:

```
ALL tests on the disk /dev/rdsk/clt1d0s2 have PASSED
```

```
The disk is now ready to be configured for I/O Fencing on node
system01
```

```
...
```

```
Removing test keys and temporary files, if any ...
```

```
.
```

```
.
```

- 7 Run the vxfcntlsthdw utility for each disk you intend to verify.

Testing the shared disks listed in a file using the vxfcntlsthdw -f option

Use the -f option to test disks that are listed in a text file. Review the following example procedure.

To test the shared disks listed in a file

- 1 Create a text file `disks_test` to test two disks shared by systems `system01` and `system02` that might resemble:

```
system01 /dev/rdisk/c2t2d1s2 system02 /dev/rdisk/c3t2d1s2
system01 /dev/rdisk/c2t2d1s2 system02 /dev/rdisk/c3t2d1s2
```

where the first disk is listed in the first line and is seen by `system01` as `/dev/rdisk/c2t2d1s2` and by `system02` as `/dev/rdisk/c3t2d1s2`. The other disk, in the second line, is seen as `/dev/rdisk/c2t2d2s2` from `system01` and `/dev/rdisk/c3t2d2s2` from `system02`. Typically, the list of disks could be extensive.

- 2 To test the disks, enter the following command:

```
# vxfentsthdw -f disks_test
```

The utility reports the test results one disk at a time, just as for the `-m` option.

Testing all the disks in a disk group using the `vxfentsthdw -g` option

Use the `-g` option to test all disks within a disk group. For example, you create a temporary disk group consisting of all disks in a disk array and test the group.

Note: Do not import the test disk group as shared; that is, do not use the `-s` option with the `vxdg import` command.

After testing, destroy the disk group and put the disks into disk groups as you need.

To test all the disks in a diskgroup

- 1 Create a diskgroup for the disks that you want to test.
- 2 Enter the following command to test the diskgroup `test_disks_dg`:

```
# vxfentsthdw -g test_disks_dg
```

The utility reports the test results one disk at a time.

Testing a disk with existing keys

If the utility detects that a coordinator disk has existing keys, you see a message that resembles:

There are Veritas I/O fencing keys on the disk. Please make sure that I/O fencing is shut down on all nodes of the cluster before continuing.

```

***** WARNING!!!!!!!!!! *****

THIS SCRIPT CAN ONLY BE USED IF THERE ARE NO OTHER ACTIVE NODES
IN THE CLUSTER!  VERIFY ALL OTHER NODES ARE POWERED OFF OR
INCAPABLE OF ACCESSING SHARED STORAGE.

If this is not the case, data corruption will result.

Do you still want to continue : [y/n] (default: n) y

```

The utility prompts you with a warning before proceeding. You may continue as long as I/O fencing is not yet configured.

About the vxfenadm utility

Administrators can use the vxfenadm command to troubleshoot and test fencing configurations.

The command's options for use by administrators are as follows:

-s	read the keys on a disk and display the keys in numeric, character, and node format
	Note: The -g and -G options are deprecated. Use the -s option.
-i	read SCSI inquiry information from device
-m	register with disks
-n	make a reservation with disks
-p	remove registrations made by other systems
-r	read reservations
-x	remove registrations

Refer to the vxfenadm(1m) manual page for a complete list of the command options.

About the I/O fencing registration key format

The keys that the vxfen driver registers on the data disks and the coordinator disks consist of eight bytes. The key format is different for the coordinator disks and data disks.

The key format of the coordinator disks is as follows:

Byte	0	1	2	3	4	5	6	7
Value	V	F	cID 0x	cID 0x	cID 0x	cID 0x	nID 0x	nID 0x

where:

- VF is the unique identifier that carves out a namespace for the keys (consumes two bytes)
- cID 0x is the LLT cluster ID in hexadecimal (consumes four bytes)
- nID 0x is the LLT node ID in hexadecimal (consumes two bytes)

The vxfen driver uses this key format in both scsi3 mode and customized mode of I/O fencing.

The key format of the data disks that are configured as failover disk groups under VCS is as follows:

Byte	0	1	2	3	4	5	6	7
Value	A+nID	V	C	S				

where nID is the LLT node ID

For example: If the node ID is 1, then the first byte has the value as B ('A' + 1 = B).

The key format of the data disks configured as parallel disk groups under CVM is as follows:

Byte	0	1	2	3	4	5	6	7
Value	A+nID	P	G	R	DGcount	DGcount	DGcount	DGcount

where DGcount is the count of disk group in the configuration (consumes four bytes).

By default, CVM uses unique fencing key for each disk group. However, some arrays have a restriction on the total number of unique keys that can be registered. In such cases, you can use the `same_key_for_alldgs` tunable parameter to change the default behavior. The default value of the parameter is off. If your configuration hits the storage array limit on total number of unique keys, you can turn the value on using the `vxdefault` command as follows:

```
# vxdefault set same_key_for_alldgs on
# vxdefault list
KEYWORD                CURRENT-VALUE    DEFAULT-VALUE
...
same_key_for_alldgs    on               off
...
```

If the tunable is changed to 'on', all subsequent keys that the CVM generates on disk group imports or creates have '0000' as their last four bytes (DGcount is 0). You must deport and re-import all the disk groups that are already imported for the changed value of the `same_key_for_alldgs` tunable to take effect.

Displaying the I/O fencing registration keys

You can display the keys that are currently assigned to the disks using the `vxfenadm` command.

The variables such as `disk_7`, `disk_8`, and `disk_9` in the following procedure represent the disk names in your setup.

To display the I/O fencing registration keys

- 1 To display the key for the disks, run the following command:

```
# vxfenadm -s disk_name
```

For example:

- To display the key for the coordinator disk `/dev/rdisk/c1t1d0s2` from the system with node ID 1, enter the following command:

```
# vxfenadm -s /dev/rdisk/c1t1d0s2
key[1]:
[Numeric Format]: 86,70,68,69,69,68,48,48
[Character Format]: VFDEED00
* [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: system01
```

The `-s` option of `vxfenadm` displays all eight bytes of a key value in three formats. In the numeric format,

- The first two bytes, represent the identifier VF, contains the ASCII value 86, 70.
- The next four bytes contain the ASCII value of the cluster ID 57069 encoded in hex (0xDEED) which are 68, 69, 69, 68.
- The remaining bytes contain the ASCII value of the node ID 0 (0x00) which are 48, 48. Node ID 1 would be 01 and node ID 10 would be 0A.

An asterisk before the Node Format indicates that the `vxfenadm` command is run from the node of a cluster where LLT is configured and running.

- To display the keys on a CVM parallel disk group:

```
# vxfenadm -s /dev/vx/rdmp/disk_7
```

```
Reading SCSI Registration Keys...
```

```
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 66,80,71,82,48,48,48,49
  [Character Format]: BPGR0001
  [Node Format]: Cluster ID: unknown Node ID: 1 Node Name: system02
```

- To display the keys on a VCS failover disk group:

```
# vxfenadm -s /dev/vx/rdmp/disk_8
```

```
Reading SCSI Registration Keys...
```

```
Device Name: /dev/vx/rdmp/disk_8
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 65,86,67,83,0,0,0,0
  [Character Format]: AVCS
  [Node Format]: Cluster ID: unknown Node ID: 0 Node Name: system01
```

- 2 To display the keys that are registered in all the disks specified in a disk file:

```
# vxfenadm -s all -f disk_filename
```

For example:

To display all the keys on coordinator disks:

```
# vxfenadm -s all -f /etc/vxfentab
```

```
Device Name: /dev/vx/rdmp/disk_9
Total Number Of Keys: 2
key[0]:
  [Numeric Format]: 86,70,66,69,65,68,48,50
  [Character Format]: VFBEAD02
  [Node Format]: Cluster ID: 48813 Node ID: 2 Node Name: unknown
key[1]:
  [Numeric Format]: 86,70,68,69,69,68,48,48
  [Character Format]: VFDEED00
* [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: system01
```

You can verify the cluster ID using the `lltstat -C` command, and the node ID using the `lltstat -N` command. For example:

```
# lltstat -C
57069
```

If the disk has keys which do not belong to a specific cluster, then the `vxfenadm` command cannot look up the node name for the node ID and hence prints the node name as unknown. For example:

```
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 1
key[0]:
[Numeric Format]: 86,70,45,45,45,45,48,49
[Character Format]: VF---01
[Node Format]: Cluster ID: unknown Node ID: 1 Node Name: system02
```

For disks with arbitrary format of keys, the `vxfenadm` command prints all the fields as unknown. For example:

```
[Numeric Format]: 65,66,67,68,49,50,51,45
[Character Format]: ABCD123-
[Node Format]: Cluster ID: unknown Node ID: unknown
Node Name: unknown
```

Verifying that the nodes see the same disk

To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option to verify that the same serial number for the LUN is returned on all paths to the LUN.

For example, an EMC disk is accessible by the `/dev/rdisk/c2t13d0s2` path on node A and the `/dev/rdisk/c2t11d0s2` path on node B.

To verify that the nodes see the same disks

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SFCFSA.
- 2 From node A, enter the following command:

```
# vxfenadm -i /dev/rdisk/c2t13d0s2
```

```
Vendor id      : EMC  
Product id    : SYMMETRIX  
Revision     : 5567  
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the /dev/rdisk/c2t11d0s2 path.

On a disk from another manufacturer, Hitachi Data Systems, the output is different and may resemble:

```
# vxfenadm -i /dev/rdisk/c2t1d0s2
```

```
Vendor id      : HITACHI  
Product id    : OPEN-3      -SUN  
Revision     : 0117  
Serial Number : 0401EB6F0002
```

Refer to the vxfenadm(1M) manual page for more information.

About the vxfenclearpre utility

You can use the vxfenclearpre utility to remove SCSI-3 registrations and reservations on the disks.

See [“Removing preexisting keys”](#) on page 404.

This utility currently does not support server-based fencing. You must manually resolve any preexisting split-brain with server-based fencing configuration.

Removing preexisting keys

If you encountered a split-brain condition, use the vxfenclearpre utility to remove SCSI-3 registrations and reservations on the coordinator disks as well as on the data disks in all shared disk groups.

You can also use this procedure to remove the registration and reservation keys created by another node from a disk.

To clear keys after split-brain

- 1 Stop VCS on all nodes.

```
# hastop -all
```

- 2 Make sure that the port h is closed on all the nodes. Run the following command on each node to verify that the port h is closed:

```
# gabconfig -a
```

Port h must not appear in the output.

- 3 Stop I/O fencing on all nodes. Enter the following command on each node:

```
# svcadm disable -t vxfen
```

- 4 If you have any applications that run outside of VCS control that have access to the shared storage, then shut down all other nodes in the cluster that have access to the shared storage. This prevents data corruption.

- 5 Start the vxfclearpre script:

```
# vxfclearpre
```

- 6 Read the script's introduction and warning. Then, you can choose to let the script run.

```
Do you still want to continue: [y/n] (default : n) y
```

In some cases, informational messages resembling the following may appear on the console of one of the nodes in the cluster when a node is ejected from a disk/LUN. You can ignore these informational messages.

```
<date> <system name> scsi: WARNING: /sbus@3,0/lpfs@0,0/  
sd@0,1(sd9l):  
<date> <system name> Error for Command: <undecoded  
cmd 0x5f> Error Level: Informational  
<date> <system name> scsi: Requested Block: 0 Error Block 0  
<date> <system name> scsi: Vendor: <vendor> Serial Number:  
0400759B006E  
<date> <system name> scsi: Sense Key: Unit Attention  
<date> <system name> scsi: ASC: 0x2a (<vendor unique code  
0x2a>), ASCQ: 0x4, FRU: 0x0
```

The script cleans up the disks and displays the following status messages.

```
Cleaning up the coordinator disks...
```

```
Cleaning up the data disks for all shared disk groups...
```

```
Successfully removed SCSI-3 persistent registration and  
reservations from the coordinator disks as well as the  
shared data disks.
```

```
You can retry starting fencing module. In order to  
restart the whole product, you might want to  
reboot the system.
```

- 7 Start the fencing module.

```
# svcadm enable vxfen
```

- 8 Start VCS on all nodes.

```
# hastart
```

About the vxfsnwap utility

The vxfsnwap utility allows you to add, remove, and replace coordinator points in a cluster that is online. The utility verifies that the serial number of the new disks are identical on all the nodes and the new disks can support I/O fencing.

This utility supports both disk-based and server-based fencing.

Refer to the vxfsnwap(1M) manual page.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for details on the I/O fencing requirements.

You can replace the coordinator disks without stopping I/O fencing in the following cases:

- The disk becomes defective or inoperable and you want to switch to a new diskgroup.
See [“Replacing I/O fencing coordinator disks when the cluster is online”](#) on page 408.
See [“Replacing the coordinator diskgroup in a cluster that is online”](#) on page 411.
If you want to replace the coordinator disks when the cluster is offline, you cannot use the vxfsnwap utility. You must manually perform the steps that the utility does to replace the coordinator disks.
- You want to switch the disk interface between raw devices and DMP devices.
See [“Changing the disk interaction policy in a cluster that is online”](#) on page 413.
- The keys that are registered on the coordinator disks are lost.
In such a case, the cluster might panic when a network partition occurs. You can replace the coordinator disks with the same disks using the vxfsnwap command. During the disk replacement, the missing keys register again without any risk of data corruption.
See [“Refreshing lost keys on coordinator disks”](#) on page 416.

In server-based fencing configuration, you can use the vxfsnwap utility to perform the following tasks:

- Perform a planned replacement of customized coordination points (CP servers or SCSI-3 disks).
See [“Replacing coordination points for server-based fencing in an online cluster”](#) on page 426.
- Refresh the I/O fencing keys that are registered on the coordination points.
See [“Refreshing registration keys on the coordination points for server-based fencing”](#) on page 428.

You can also use the `vxftswap` utility to migrate between the disk-based and the server-based fencing without incurring application downtime in the SFCFSHA cluster.

See [“Migrating from disk-based to server-based fencing in an online cluster”](#) on page 438.

See [“Migrating from server-based to disk-based fencing in an online cluster”](#) on page 443.

If the `vxftswap` operation is unsuccessful, then you can use the `-a cancel` of the `vxftswap` command to manually roll back the changes that the `vxftswap` utility does.

- For disk-based fencing, use the `vxftswap -g diskgroup -a cancel` command to cancel the `vxftswap` operation.
You must run this command if a node fails during the process of disk replacement, or if you aborted the disk replacement.
- For server-based fencing, use the `vxftswap -a cancel` command to cancel the `vxftswap` operation.

Replacing I/O fencing coordinator disks when the cluster is online

Review the procedures to add, remove, or replace one or more coordinator disks in a cluster that is operational.

Warning: The cluster might panic if any node leaves the cluster membership before the `vxftswap` script replaces the set of coordinator disks.

To replace a disk in a coordinator diskgroup when the cluster is online

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfenadm -d

I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
    1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Import the coordinator disk group.

The file `/etc/vxfendg` includes the name of the disk group (typically, `vxfencoordg`) that contains the coordinator disks, so use the command:

```
# vxdg -tFC import `cat /etc/vxfendg`
```

where:

-t specifies that the disk group is imported only until the node restarts.

-f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.

-C specifies that any import locks are removed.

- 4 If your setup uses VRTSvxvm version 5.1 RP2 (or later) or 5.1 SP1 (or later), then skip to step 5. You need not set `coordinator=off` to add or remove disks. For other VxVM versions, perform this step:

Turn off the coordinator attribute value for the coordinator disk group.

```
# vxdg -g vxfencoordg set coordinator=off
```

- 5 To remove disks from the coordinator disk group, use the VxVM disk administrator utility `vxdiskadm`.
- 6 Perform the following steps to add new disks to the coordinator disk group:
 - Add new disks to the node.

- Initialize the new disks as VxVM disks.
- Check the disks for I/O fencing compliance.
- Add the new disks to the coordinator disk group and set the coordinator attribute value as "on" for the coordinator disk group.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for detailed instructions.

Note that though the disk group content changes, the I/O fencing remains in the same state.

- 7 From one node, start the `vxfsnwap` utility. You must specify the diskgroup to the utility.

The utility performs the following tasks:

- Backs up the existing `/etc/vxfentab` file.
- Creates a test file `/etc/vxfentab.test` for the diskgroup that is modified on each node.
- Reads the diskgroup you specified in the `vxfsnwap` command and adds the diskgroup to the `/etc/vxfentab.test` file on each node.
- Verifies that the serial number of the new disks are identical on all the nodes. The script terminates if the check fails.
- Verifies that the new disks can support I/O fencing on each node.

- 8 If the disk verification passes, the utility reports success and asks if you want to commit the new set of coordinator disks.

- 9 Confirm whether you want to clear the keys on the coordination points and proceed with the `vxfsnwap` operation.

```
Do you want to clear the keys on the coordination points
and proceed with the vxfsnwap operation? [y/n] (default: n) y
```

- 10 Review the message that the utility displays and confirm that you want to commit the new set of coordinator disks. Else skip to step 11.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfentab.test` file to the `/etc/vxfentab` file.

- 11 If you do not want to commit the new set of coordinator disks, answer `n`.

The `vxfsnwap` utility rolls back the disk replacement operation.

Replacing the coordinator diskgroup in a cluster that is online

You can also replace the coordinator diskgroup using the `vxftswap` utility. The following example replaces the coordinator disk group `vxftcoorddg` with a new disk group `vxftendg`.

To replace the coordinator diskgroup

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxftadm -d
```

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: dmp
```

```
Cluster Members:
```

```
  * 0 (system01)
```

```
  1 (system02)
```

```
RFSM State Information:
```

```
  node 0 in state 8 (running)
```

```
  node 1 in state 8 (running)
```

- 3 Find the name of the current coordinator diskgroup (typically `vxftcoorddg`) that is in the `/etc/vxftendg` file.

```
# cat /etc/vxftendg
```

```
vxftcoorddg
```

- 4 Find the alternative disk groups available to replace the current coordinator diskgroup.

```
# vxftdisk -o alldgs list
```

- 5 Validate the new disk group for I/O fencing compliance. Run the following command:

```
# vxftsthdw -c vxftendg
```

See [“Testing the coordinator disk group using vxftsthdw -c option”](#) on page 394.

- 6 If the new disk group is not already deported, run the following command to deport the disk group:

```
# vxdg deport vxfendg
```

- 7 Make sure that the `/etc/vxfenmode` file is updated to specify the correct disk policy.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more information.

- 8 From any node, start the `vxfsnwap` utility. For example, if `vxfendg` is the new diskgroup that you want to use as the coordinator diskgroup:

```
# vxfsnwap -g vxfendg [-n]
```

The utility performs the following tasks:

- Backs up the existing `/etc/vxfentab` file.
 - Creates a test file `/etc/vxfentab.test` for the diskgroup that is modified on each node.
 - Reads the diskgroup you specified in the `vxfsnwap` command and adds the diskgroup to the `/etc/vxfentab.test` file on each node.
 - Verifies that the serial number of the new disks are identical on all the nodes. The script terminates if the check fails.
 - Verifies that the new disk group can support I/O fencing on each node.
- 9 If the disk verification passes, the utility reports success and asks if you want to replace the coordinator disk group.
- 10 Confirm whether you want to clear the keys on the coordination points and proceed with the `vxfsnwap` operation.

```
Do you want to clear the keys on the coordination points  
and proceed with the vxfsnwap operation? [y/n] (default: n) y
```

- 11 Review the message that the utility displays and confirm that you want to replace the coordinator disk group. Else skip to step 14.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfentab.test` file to the `/etc/vxfentab` file.

The utility also updates the `/etc/vxfendg` file with this new diskgroup.

- 12 Set the coordinator attribute value as "on" for the new coordinator disk group.

```
# vxdg -g vxfsendg set coordinator=on
```

Set the coordinator attribute value as "off" for the old disk group.

```
# vxdg -g vxfsencoorddg set coordinator=off
```

- 13 Verify that the coordinator disk group has changed.

```
# cat /etc/vxfsendg
vxfsendg
```

The swap operation for the coordinator disk group is complete now.

- 14 If you do not want to replace the coordinator disk group, answer n at the prompt.

The vxfsenswap utility rolls back any changes to the coordinator diskgroup.

Changing the disk interaction policy in a cluster that is online

In a cluster that is online, you can change the disk interaction policy from dmp to raw using the vxfsenswap utility.

To change the disk interaction policy

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfsenadm -d
```

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: dmp
```

```
Cluster Members:
```

```
* 0 (system01)
```

```
1 (system02)
```

```
RFSM State Information:
```

```
node 0 in state 8 (running)
```

```
node 1 in state 8 (running)
```

- 3 On each node in the cluster, edit the `/etc/vxfenmode` file to change the disk policy.

```
# cat /etc/vxfenmode
vxfen_mode=scsi3
scsi3_disk_policy=raw
```

- 4 From any node, start the `vxfenswap` utility:

```
# vxfenswap -g vxfencoordg [-n]
```

- 5 Verify the change in the disk policy.

```
# vxfenadm -d
```

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: raw
```

Adding disks from a recovered site to the coordinator diskgroup

In a campus cluster environment, consider a case where the primary site goes down and the secondary site comes online with a limited set of disks. When the primary site restores, the primary site's disks are also available to act as coordinator disks. You can use the `vxfenswap` utility to add these disks to the coordinator diskgroup.

To add new disks from a recovered site to the coordinator diskgroup

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfenadm -d
```

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: dmp
```

```
Cluster Members:
```

```
  * 0 (system01)
```

```
  1 (system02)
```

```
RFSM State Information:
```

```
  node 0 in state 8 (running)
```

```
  node 1 in state 8 (running)
```

- 3 Verify the name of the coordinator diskgroup.

```
# cat /etc/vxfendg
```

```
vxfencoorddg
```

- 4 Run the following command:

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
clt1d0s2	auto:cdsdisk	-	(vxfencoorddg)	online
c2t1d0s2	auto	-	-	offline
c3t1d0s2	auto	-	-	offline

- 5 Verify the number of disks used in the coordinator diskgroup.

```
# vxfenconfig -l
```

```
I/O Fencing Configuration Information:
```

```
=====
```

```
Count : 1
```

```
Disk List
```

Disk Name	Major	Minor	Serial Number	Policy
/dev/vx/rdmp/clt1d0s2	32	48	R450 00013154	0312 dmp

- When the primary site comes online, start the vxfenswap utility on any node in the cluster:

```
# vxfenswap -g vxfencoorddg [-n]
```

- Verify the count of the coordinator disks.

```
# vxfenconfig -l
I/O Fencing Configuration Information:
=====
Single Disk Flag      : 0
Count                 : 3
Disk List
Disk Name             Major  Minor  Serial Number      Policy
/dev/vx/rdmp/clt1d0s2      32   48   R450 00013154 0312      dmp
/dev/vx/rdmp/c2t1d0s2      32   32   R450 00013154 0313      dmp
/dev/vx/rdmp/c3t1d0s2      32   16   R450 00013154 0314      dmp
```

Refreshing lost keys on coordinator disks

If the coordinator disks lose the keys that are registered, the cluster might panic when a network partition occurs.

You can use the vxfenswap utility to replace the coordinator disks with the same disks. The vxfenswap utility registers the missing keys during the disk replacement.

To refresh lost keys on coordinator disks

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfenadm -d
```

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Run the following command to view the coordinator disks that do not have keys:

```
# vxfenadm -s all -f /etc/vxfentab
```

```
Device Name: /dev/vx/rdmp/cltld0s2
Total Number of Keys: 0
No keys...
...
```

- 4 On any node, run the following command to start the vxfenswap utility:

```
# vxfenswap -g vxfencoorddg [-n]
```

- 5 Verify that the keys are atomically placed on the coordinator disks.

```
# vxfenadm -s all -f /etc/vxfentab
```

```
Device Name: /dev/vx/rdmp/cltld0s2
Total Number of Keys: 4
...
```

About administering the coordination point server

This section describes how to perform administrative and maintenance tasks on the coordination point server (CP server).

For more information about the `cpsadm` command and the associated command options, see the `cpsadm(1M)` manual page.

Environment variables associated with the coordination point server

[Table 15-4](#) describes the environment variables that are required for the `cpsadm` command. The `cpsadm` command detects these environment variables and uses their value when communicating with the CP server. They are used to authenticate and authorize the user.

Note: The environment variables are not required when the `cpsadm` command is run on the CP server. The environment variables are required when the `cpsadm` command is run on the SFCFSA cluster nodes.

Table 15-4 `cpsadm` command environment variables

Environment variable	Description
CPS_USERNAME	This is the fully qualified username as configured in VxSS (vssat showcred).
CPS_DOMAINTYPE	One of the following values: ■ vx ■ unixpwd ■ nis ■ nisplus ■ ldap

The environment variables must be exported directly on the shell before running the `cpsadm` command. For example,

```
# export CPS_USERNAME=  
# export CPS_DOMAINTYPE=
```

Additionally, the username and domaintype values are the same as those added onto the CP server. To view these values run the following `cpsadm` command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp_server* is the CP server's virtual IP address or virtual hostname.

CP server operations (cpsadm)

[Table 15-5](#) lists coordination point server (CP server) operations and required privileges.

Table 15-5 User privileges for CP server operations

CP server operations	CP server Operator	CP server Admin
add_cluster	—	✓
rm_clus	—	✓
add_node	✓	✓
rm_node	✓	✓
add_user	—	✓
rm_user	—	✓
add_clus_to_user	—	✓
rm_clus_from_user	—	✓
reg_node	✓	✓
unreg_node	✓	✓
preempt_node	✓	✓
list_membership	✓	✓
list_nodes	✓	✓
list_users	✓	✓
halt_cps	—	✓
db_snapshot	—	✓
ping_cps	✓	✓
client_preupgrade	✓	✓
server_preupgrade	✓	✓
list_protocols	✓	✓
list_version	✓	✓

Table 15-5 User privileges for CP server operations (*continued*)

CP server operations	CP server Operator	CP server Admin
list_ports	–	✓
add_port	–	✓
rm_port	–	✓

Adding and removing SFCFSHA cluster entries from the CP server database

- To add a SFCFSHA cluster to the CP server database

Type the following command:

```
# cpsadm -s cp_server -a add_clus -c cluster_name -u uuid
```

- To remove a SFCFSHA cluster from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a rm_clus -u uuid
```

cp_server The CP server's virtual IP address or virtual hostname.

cluster_name The SFCFSHA cluster name.

uuid The UUID (Universally Unique ID) of the SFCFSHA cluster.

Adding and removing a SFCFSHA cluster node from the CP server database

- To add a SFCFSHA cluster node from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a add_node -u uuid -n nodeid
-h host
```

- To remove a SFCFSHA cluster node from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a rm_node -u uuid -n nodeid
```

cp_server The CP server's virtual IP address or virtual hostname.

<i>uuid</i>	The UUID (Universally Unique ID) of the SFCFSHA cluster.
<i>nodeid</i>	The node id of the SFCFSHA cluster node.
<i>host</i>	Hostname

Adding or removing CP server users

- To add a user

Type the following command:

```
# cpsadm -s cp_server -a add_user -e user_name -f user_role
-g domain_type -u uuid
```

- To remove a user

Type the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

<i>cp_server</i>	The CP server's virtual IP address or virtual hostname.
<i>user_name</i>	The user to be added to the CP server configuration.
<i>user_role</i>	The user role, either cps_admin or cps_operator.
<i>domain_type</i>	The domain type, for example vx, unixpwd, nis, etc.
<i>uuid</i>	The UUID (Universally Unique ID) of the SFCFSHA cluster.

Listing the CP server users

To list the CP server users

Type the following command:

```
# cpsadm -s cp_server -a list_users
```

Listing the nodes in all the SFCFSHA clusters

To list the nodes in all the SFCFSHA cluster

Type the following command:

```
# cpsadm -s cp_server -a list_nodes
```

Listing the membership of nodes in the SFCFSHA cluster

To list the membership of nodes in SFCFSHA cluster

Type the following command:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

cp_server The CP server's virtual IP address or virtual hostname.

cluster_name The SFCFSHA cluster name.

Preempting a node

Use the following command to preempt a node.

To preempt a node

◆ Type the following command:

```
# cpsadm -s cp_server -a preempt_node -u uuid -n nodeid  
-v victim_node id
```

cp_server The CP server's virtual IP address or virtual hostname.

uuid The UUID (Universally Unique ID) of the SFCFSHA cluster.

nodeid The node id of the SFCFSHA cluster node.

victim_node id Node id of one or more victim nodes.

Registering and unregistering a node

■ To register a node

Type the following command:

```
# cpsadm -s cp_server -a reg_node -u uuid -n nodeid
```

■ To unregister a node

Type the following command:

```
# cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
```

cp_server The CP server's virtual IP address or virtual hostname.

uuid The UUID (Universally Unique ID) of the SFCFSHA cluster.

nodeid The nodeid of the SFCFSHA cluster node.

Enable and disable access for a user to a SFCFSHA cluster

- To enable access for a user to a SFCFSHA cluster

Type the following command:

```
# cpsadm -s cp_server -a add_clus_to_user -e user
-f user_role -g domain_type -u uuid
```

- To disable access for a user to a SFCFSHA cluster

Type the following command:

```
# cpsadm -s cp_server -a rm_clus_from_user -e user_name
-f user_role -g domain_type -u uuid
```

cp_server The CP server's virtual IP address or virtual hostname.

user_name The user name to be added to the CP server.

user_role The user role, either cps_admin or cps_operator.

domain_type The domain type, for example vx, unixpwd, nis, etc.

uuid The UUID (Universally Unique ID) of the SFCFSHA cluster

Starting and stopping CP server outside VCS control

You can start or stop coordination point server (CP server) outside VCS control.

To start CP server outside VCS control

- 1 Run the `vxcpserv` binary directly:

```
# /opt/VRTScps/bin/vxcpserv
```

If the command is successful, the command immediately returns without any message.

- 2 Verify the log file `/var/VRTScps/log/cpserver_A.log` to confirm the state of the CP server.

To stop CP server outside VCS control

- 1 Run the following command:

```
# cpsadm -s cp_server -a halt_cps
```

The variable *cp_server* represents the CP server's virtual IP address or virtual host name and *port_number* represents the port number on which the CP server is listening.

- 2 Verify the log file `/var/VRTScps/log/cpsserver_A.log` to confirm that the CP server received the halt message and has shut down.

Checking the connectivity of CP servers

To check the connectivity of a CP server

Type the following command:

```
# cpsadm -s cp_server -a ping_cps
```

Adding and removing virtual IP addresses and ports for CP server at run-time

You can use more than one virtual IP address for coordination point server (CP server) communication. You can assign port numbers to each of the virtual IP addresses.

You can use the `cpsadm` command if you want to add or remove virtual IP addresses and ports after your initial CP server setup. However, these virtual IP addresses and ports that you add or remove does not change the `vxcps.conf` file. So, these changes do not persist across CP server restarts.

See the `cpsadm(1m)` manual page for more details.

To add and remove virtual IP addresses and ports for CP servers at run-time

- 1 To list all the ports that the CP server is configured to listen on, run the following command:

```
# cpsadm -s cp_server -a list_ports
```

If the CP server has not been able to successfully listen on a given port at least once, then the Connect History in the output shows never. If the IP addresses are down when the vxcperv process starts, vxcperv binds to the IP addresses when the addresses come up later. For example:

```
# cpsadm -s 127.0.0.1 -a list_ports
```

IP Address	Connect History
[10.209.79.60]:14250	once
[10.209.79.61]:56789	once
[10.209.78.252]:14250	never
[192.10.10.32]:14250	once

CP server does not actively monitor port health. If the CP server successfully listens on any IP:port at least once, then the Connect History for that IP:port shows once even if the port goes down later during CP server's lifetime. You can obtain the latest status of the IP address from the corresponding IP resource state that is configured under VCS.

- 2 To add a new port (IP:port) for the CP server without restarting the CP server, run the following command:

```
# cpsadm -s cp_server -a add_port  
-i ip_address -r port_number
```

For example:

```
# cpsadm -s 127.0.0.1 -a add_port -i 10.209.78.52 -r 14250  
Port [10.209.78.52]:14250 successfully added.
```

- 3 To stop the CP server from listening on a port (IP:port) without restarting the CP server, run the following command:

```
# cpsadm -s cp_server -a rm_port  
-i ip_address -r port_number
```

For example:

```
# cpsadm -s 10.209.78.52 -a rm_port -i 10.209.78.252  
No port specified. Assuming default port i.e 14250  
Port [10.209.78.252]:14250 successfully removed.
```

Taking a CP server database snapshot

To take a CP server database snapshot

Type the following command:

```
# cpsadm -s cp_server -a db_snapshot
```

Replacing coordination points for server-based fencing in an online cluster

Use the following procedure to perform a planned replacement of customized coordination points (CP servers or SCSI-3 disks) without incurring application downtime on an online SFCFSHA cluster.

Note: If multiple clusters share the same CP server, you must perform this replacement procedure in each cluster.

You can use the `vxfsnwap` utility to replace coordination points when fencing is running in customized mode in an online cluster, with `vxfen_mechanism=cps`. The utility also supports migration from server-based fencing (`vxfen_mode=customized`) to disk-based fencing (`vxfen_mode=scsi3`) and vice versa in an online cluster.

However, if the SFCFSHA cluster has fencing disabled (`vxfen_mode=disabled`), then you must take the cluster offline to configure disk-based or server-based fencing.

See [“Deployment and migration scenarios for CP server”](#) on page 430.

You can cancel the coordination point replacement operation at any time using the `vxfsnwap -a cancel` command.

See [“About the vxfsnwap utility”](#) on page 407.

To replace coordination points for an online cluster

- 1 Ensure that the SFCFSHA cluster nodes and users have been added to the new CP server(s). Run the following commands:

```
# cpsadm -s cpserver -a list_nodes
# cpsadm -s cpserver -a list_users
```

If the SFCFSHA cluster nodes are not present here, prepare the new CP server(s) for use by the SFCFSHA cluster.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for instructions.

- 2 Ensure that fencing is running on the cluster using the old set of coordination points and in customized mode.

For example, enter the following command:

```
# vxfenadm -d
```

The command returns:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: <version>
Fencing Mode: Customized
Cluster Members:
* 0 (galaxy)
1 (nebula)
RFSM State Information:
node 0 in state 8 (running)
node 1 in state 8 (running)
```

- 3 Create a new `/etc/vxfenmode.test` file on each SFCFSHA cluster node with the fencing configuration changes such as the CP server information.

Review and if necessary, update the `vxfenmode` parameters for security, the coordination points, and if applicable to your configuration, `vxfendg`.

Refer to the text information within the `vxfenmode` file for additional information about these parameters and their new possible values.

- 4 From one of the nodes of the cluster, run the `vxfenswap` utility.

The `vxfenswap` utility requires secure ssh connection to all the cluster nodes. Use `-n` to use rsh instead of default ssh.

```
# vxfenswap [-n]
```

- 5 Review the message that the utility displays and confirm whether you want to commit the change.

- If you do not want to commit the new fencing configuration changes, press Enter or answer n at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) n
```

The `vxferswap` utility rolls back the migration operation.

- If you want to commit the new fencing configuration changes, answer y at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully completes the operation, the utility moves the `/etc/vxfermode.test` file to the `/etc/vxfermode` file.

- 6 Confirm the successful execution of the `vxferswap` utility by checking the coordination points currently used by the `vxfer` driver.

For example, run the following command:

```
# vxferconfig -l
```

Refreshing registration keys on the coordination points for server-based fencing

Replacing keys on a coordination point (CP server) when the SFCFSHA cluster is online involves refreshing that coordination point's registrations. You can perform a planned refresh of registrations on a CP server without incurring application downtime on the SFCFSHA cluster. You must refresh registrations on a CP server if the CP server agent issues an alert on the loss of such registrations on the CP server database.

The following procedure describes how to refresh the coordination point registrations.

To refresh the registration keys on the coordination points for server-based fencing

- 1 Ensure that the SFCFSHA cluster nodes and users have been added to the new CP server(s). Run the following commands:

```
# cpsadm -s cp_server -a list_nodes
```

```
# cpsadm -s cp_server -a list_users
```

If the SFCFSHA cluster nodes are not present here, prepare the new CP server(s) for use by the SFCFSHA cluster.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for instructions.

- 2 Ensure that fencing is running on the cluster in customized mode using the coordination points mentioned in the `/etc/vxfenmode` file.

For example, enter the following command:

```
# vxfenadm -d
```

```
=====
Fencing Protocol Version: 201
Fencing Mode: CUSTOMIZED
Cluster Members:
* 0 (galaxy)
1 (nebula)
RFSM State Information:
node 0 in state 8 (running)
node 1 in state 8 (running)
```

- 3 List the coordination points currently used by I/O fencing :

```
# vxfenconfig -l
```

- 4 Run the `vxfsenwap` utility from one of the nodes of the cluster.

The `vxfsenwap` utility requires secure ssh connection to all the cluster nodes. Use `-n` to use rsh instead of default ssh.

For example:

```
# vxfsenwap [-n]
```

The command returns:

```
VERITAS vxfsenwap version <version> <platform>
The logfile generated for vxfsenwap is
/var/VRTSvcs/log/vxfen/vxfsenwap.log.
19156
Please Wait...
VXFEN vxfsenconfig NOTICE Driver will use customized fencing
- mechanism cps
Validation of coordination points change has succeeded on
all nodes.
You may commit the changes now.
WARNING: This may cause the whole cluster to panic
if a node leaves membership before the change is complete.
```

- 5 You are then prompted to commit the change. Enter **y** for yes.

The command returns a confirmation of successful coordination point replacement.

- 6 Confirm the successful execution of the `vxfsenwap` utility. If CP agent is configured, it should report ONLINE as it succeeds to find the registrations on coordination points. The registrations on the CP server and coordinator disks can be viewed using the `cpsadm` and `vxfsenadm` utilities respectively.

Note that a running online coordination point refreshment operation can be canceled at any time using the command:

```
# vxfsenwap -a cancel
```

Deployment and migration scenarios for CP server

[Table 15-6](#) describes the supported deployment and migration scenarios, and the procedures you must perform on the SFCFSA cluster and the CP server.

Table 15-6 CP server deployment and migration scenarios

Scenario	CP server	SFCFSHA cluster	Action required
Setup of CP server for a SFCFSHA cluster for the first time	New CP server	New SFCFSHA cluster using CP server as coordination point	On the designated CP server, perform the following tasks: 1 Prepare to configure the new CP server. 2 Configure the new CP server. 3 Prepare the new CP server for use by the SFCFSHA cluster. On the SFCFSHA cluster nodes, configure server-based I/O fencing. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.
Add a new SFCFSHA cluster to an existing and operational CP server	Existing and operational CP server	New SFCFSHA cluster	On the SFCFSHA cluster nodes, configure server-based I/O fencing. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.
Replace the coordination point from an existing CP server to a new CP server	New CP server	Existing SFCFSHA cluster using CP server as coordination point	On the designated CP server, perform the following tasks: 1 Prepare to configure the new CP server. 2 Configure the new CP server. 3 Prepare the new CP server for use by the SFCFSHA cluster. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures. On a node in the SFCFSHA cluster, run the <code>vxfsnwap</code> command to move to replace the CP server: See “Replacing coordination points for server-based fencing in an online cluster” on page 426.

Table 15-6 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFSHA cluster	Action required
Replace the coordination point from an existing CP server to an operational CP server coordination point	Operational CP server	Existing SFCFSHA cluster using CP server as coordination point	On the designated CP server, prepare to configure the new CP server manually. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures. On a node in the SFCFSHA cluster, run the <code>vxfsnwap</code> command to move to replace the CP server: See "Replacing coordination points for server-based fencing in an online cluster" on page 426.
Enabling fencing in a SFCFSHA cluster with a new CP server coordination point	New CP server	Existing SFCFSHA cluster with fencing configured in disabled mode	Note: Migrating from fencing in disabled mode to customized mode incurs application downtime on the SFCFSHA cluster. On the designated CP server, perform the following tasks: 1 Prepare to configure the new CP server. 2 Configure the new CP server 3 Prepare the new CP server for use by the SFCFSHA cluster See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures. On the SFCFSHA cluster nodes, perform the following: 1 Stop all applications, VCS, and fencing on the SFCFSHA cluster. 2 To stop VCS, use the following command (to be run on all the SFCFSHA cluster nodes): <pre># hstop -local</pre> 3 Stop fencing using the following command: <pre># svcadm disable -t vxfen</pre> 4 Reconfigure I/O fencing on the SFCFSHA cluster. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.

Table 15-6 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFSHA cluster	Action required
Enabling fencing in a SFCFSHA cluster with an operational CP server coordination point	Operational CP server	Existing SFCFSHA cluster with fencing configured in disabled mode	Note: Migrating from fencing in disabled mode to customized mode incurs application downtime. On the designated CP server, prepare to configure the new CP server. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for this procedure. On the SFCFSHA cluster nodes, perform the following tasks: 1 Stop all applications, VCS, and fencing on the SFCFSHA cluster. 2 To stop VCS, use the following command (to be run on all the SFCFSHA cluster nodes): # hastop -local 3 Stop fencing using the following command: # svcadm disable -t vxfen 4 Reconfigure fencing on the SFCFSHA cluster. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.

Table 15-6 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFSHA cluster	Action required
Enabling fencing in a SFCFSHA cluster with a new CP server coordination point	New CP server	Existing SFCFSHA cluster with fencing configured in scsi3 mode	On the designated CP server, perform the following tasks: 1 Prepare to configure the new CP server. 2 Configure the new CP server 3 Prepare the new CP server for use by the SFCFSHA cluster See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures. Based on whether the cluster is online or offline, perform the following procedures: For a cluster that is online, perform the following task on the SFCFSHA cluster: ◆ Run the <code>vxfsnswap</code> command to migrate from disk-based fencing to the server-based fencing. See “Migrating from disk-based to server-based fencing in an online cluster” on page 438. For a cluster that is offline, perform the following tasks on the SFCFSHA cluster: 1 Stop all applications, VCS, and fencing on the SFCFSHA cluster. 2 To stop VCS, use the following command (to be run on all the SFCFSHA cluster nodes): # <code>hastop -local</code> 3 Stop fencing using the following command: # <code>svcadm disable -t vxfen</code> 4 Reconfigure I/O fencing on the SFCFSHA cluster. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.

Table 15-6 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFSHA cluster	Action required
Enabling fencing in a SFCFSHA cluster with an operational CP server coordination point	Operational CP server	Existing SFCFSHA cluster with fencing configured in disabled mode	On the designated CP server, prepare to configure the new CP server. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for this procedure. Based on whether the cluster is online or offline, perform the following procedures: For a cluster that is online, perform the following task on the SFCFSHA cluster: ◆ Run the <code>vx fenceswap</code> command to migrate from disk-based fencing to the server-based fencing. See "Migrating from disk-based to server-based fencing in an online cluster" on page 438. For a cluster that is offline, perform the following tasks on the SFCFSHA cluster: 1 Stop all applications, VCS, and fencing on the SFCFSHA cluster. 2 To stop VCS, use the following command (to be run on all the SFCFSHA cluster nodes): <pre># hstop -local</pre> 3 Stop fencing using the following command: <pre># svcadm disable -t vxfen</pre> 4 Reconfigure fencing on the SFCFSHA cluster. See the <i>Veritas Storage Foundation Cluster File System High Availability Installation Guide</i> for the procedures.

Table 15-6 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFSHA cluster	Action required
Refreshing registrations of SFCFSHA cluster nodes on coordination points (CP servers/ coordinator disks) without incurring application downtime	Operational CP server	Existing SFCFSHA cluster using the CP server as coordination point	On the SFCFSHA cluster run the <code>vxfsenwap</code> command to refresh the keys on the CP server: See “Refreshing registration keys on the coordination points for server-based fencing” on page 428.

Migrating from non-secure to secure setup for CP server and SFCFSHA cluster communication

The following procedure describes how to migrate from a non-secure to secure set up for the coordination point server (CP server) and SFCFSHA cluster.

To migrate from non-secure to secure setup for CP server and SFCFSHA cluster

- 1 Stop VCS on all cluster nodes that use the CP servers.

```
# hstop -all
```

- 2 Stop fencing on all the SFCFSHA cluster nodes of all the clusters.

```
# svcadm disable -t vxfen
```

- 3 Stop all the CP servers using the following command on each CP server:

```
# hagrps -offline CPSSG -any
```

- 4 Ensure that security is configured for communication on CP Servers as well as all the clients.

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for more information.

- 5
 - If CP server is hosted on an SFHA cluster, perform this step on each CP server.

Bring the mount resource in the CPSSG service group online.

```
# hares -online cpsmount -sys local_system_name
```


Complete the remaining steps.

- If CP server is hosted on a single-node VCS cluster, skip to step 8 and complete the remaining steps.
- 6 After the mount resource comes online, move the `credentials` directory from the default location to shared storage.

```
# mv /var/VRTSvcs/vcsauth/data/CPSERVER /etc/VRTSvcs/db/
```

- 7 Create softlinks on all the nodes of the CP servers.

```
# ln -s /etc/VRTScps/db/CPSERVER \  
/var/VRTSvcs/vcsauth/data/CPSERVER
```

- 8 Edit `/etc/vxcps.conf` on each CP server to set `security=1`.

- 9 Start CP servers by using the following command:

```
# hagrps -online CPSSG -any
```

- 10 Edit `/etc/VRTSvcs/conf/config/main.cf` on the first node of the cluster and remove the `UseFence=SCSI3` attribute.

Start VCS on the first node and then on all other nodes of the cluster.

- 11 Reconfigure fencing on each cluster by using the installer.

```
# /opt/VRTS/install/installsfcfsha -fencing
```

About migrating between disk-based and server-based fencing configurations

You can migrate between fencing configurations without incurring application downtime in the SFCFSHA clusters.

You can migrate from disk-based fencing to server-based fencing in the following cases:

- You want to leverage the benefits of server-based fencing.
- You want to replace faulty coordinator disks with coordination point servers (CP servers).

See [“Migrating from disk-based to server-based fencing in an online cluster”](#) on page 438.

Similarly, you can migrate from server-based fencing to disk-based fencing when you want to perform maintenance tasks on the CP server systems.

See [“Migrating from server-based to disk-based fencing in an online cluster”](#) on page 443.

Migrating from disk-based to server-based fencing in an online cluster

You can either use the installer or manually migrate from disk-based fencing to server-based fencing without incurring application downtime in the SFCFSA clusters.

See [“About migrating between disk-based and server-based fencing configurations”](#) on page 437.

You can also use response files to migrate between fencing configurations.

See [“Migrating between fencing configurations using response files”](#) on page 448.

Warning: The cluster might panic if any node leaves the cluster membership before the coordination points migration operation completes.

This section covers the following procedures:

Migrating using the script-based installer	See “To migrate from disk-based fencing to server-based fencing using the installer” on page 439.
Migrating manually	See “To manually migrate from disk-based fencing to server-based fencing” on page 442.

To migrate from disk-based fencing to server-based fencing using the installer

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the SFCFSHA cluster is online and uses disk-based fencing.

```
# vxfsadm -d
```

For example, if SFCFSHA cluster uses disk-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 On any node in the cluster, start the `installsfcfsha` with the `-fencing` option.

```
# /opt/VRTS/install/installsfcfsha -fencing
```

The `installsfcfsha` starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 4 Confirm that you want to proceed with the I/O fencing configuration.
The installer verifies whether I/O fencing is configured in enabled mode.
- 5 Confirm that you want to reconfigure I/O fencing.
- 6 Review the I/O fencing configuration options that the program presents. Type **4** to migrate to server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-4,q] 4
```

- 7 From the list of coordination points that the installer presents, select the coordination points that you want to replace.

For example:

```
Select the coordination points you would like to remove
from the currently configured coordination points:
```

- 1) emc_clariion0_62
- 2) emc_clariion0_65
- 3) emc_clariion0_66
- 4) All
- 5) None
- b) Back to previous menu

```
Enter the options separated by spaces: [1-5,b,q,?] (5)? 1 2
```

If you want to migrate to server-based fencing with no coordinator disks, type **4** to remove all the coordinator disks.

- 8 Enter the total number of new coordination points.

If you want to migrate to server-based fencing configuration with a mix of coordination points, the number you enter at this prompt must be a total of both the new CP servers and the new coordinator disks.

- 9 Enter the total number of new coordinator disks.

If you want to migrate to server-based fencing with no coordinator disks, type **0** at this prompt.

- 10 Enter the total number of virtual IP addresses or host names of the virtual IP address for each of the CP servers.
- 11 Enter the virtual IP addresses or host names of the virtual IP address for each of the CP servers.

- 12 Verify and confirm the coordination points information for the fencing reconfiguration.

- 13 Review the output as the installer performs the following tasks:

- Removes the coordinator disks from the coordinator disk group.
- Updates the application cluster details on each of the new CP servers.
- Prepares the `vxfenmode.test` file on all nodes.
- Runs the `vxfsnswap` script.
Note the location of the `vxfsnswap.log` file which you can access in the event of any problem with the configuration process.
- Completes the I/O fencing migration.

- 14** If you want to send this installation information to Symantec, answer **y** at the prompt.

```
Would you like to send the information about this installation  
to Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

- 15** After the migration is complete, verify the change in the fencing mode.

```
# vxfenadm -d
```

For example, after the migration from disk-based fencing to server-based fencing in the SFCFSHA cluster:

```
I/O Fencing Cluster Information:  
=====
```

Fencing Protocol Version:	201
Fencing Mode:	Customized
Fencing Mechanism:	cps
Cluster Members:	
* 0	(system01)
1	(system02)
RFSM State Information:	
node 0	in state 8 (running)
node 1	in state 8 (running)

- 16** Verify the current coordination points that the vxfen driver uses.

```
# vxfenconfig -l
```

To manually migrate from disk-based fencing to server-based fencing

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the SFCFSHA cluster is online and uses disk-based fencing.

```
# vxfenadm -d
```

For example, if SFCFSHA cluster uses disk-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Make sure that you performed the following tasks on the designated CP server:
 - Preparing to configure the new CP server.
 - Configuring the new CP server
 - Preparing the new CP server for use by the SFCFSHA cluster

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for the procedures.

- 4 Create a new `/etc/vxfenmode.test` file on each SFCFSHA cluster node with the fencing configuration changes such as the CP server information.

Refer to the sample `vxfenmode` files in the `/etc/vxfen.d` folder.

- 5 From any node in the SFCFSHA cluster, start the `vxferswap` utility:

```
# vxferswap [-n]
```

- 6 Review the message that the utility displays and confirm whether you want to commit the change.

- If you do not want to commit the new fencing configuration changes, press Enter or answer **n** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) n
```

The `vxfsnwap` utility rolls back the migration operation.

- If you want to commit the new fencing configuration changes, answer **y** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfenmode.test` file to the `/etc/vxfenmode` file.

- 7 After the migration is complete, verify the change in the fencing mode.

```
# vxfsnadm -d
```

For example, after the migration from disk-based fencing to server-based fencing in the SFCFSHA cluster:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 8 Verify the current coordination points that the `vxfsn` driver uses.

```
# vxfsnconfig -l
```

Migrating from server-based to disk-based fencing in an online cluster

You can either use the installer or manually migrate from server-based fencing to disk-based fencing without incurring application downtime in the SFCFSHA clusters.

See [“About migrating between disk-based and server-based fencing configurations”](#) on page 437.

You can also use response files to migrate between fencing configurations.

See [“Migrating between fencing configurations using response files”](#) on page 448.

Warning: The cluster might panic if any node leaves the cluster membership before the coordination points migration operation completes.

This section covers the following procedures:

Migrating using the script-based installer	See “To migrate from server-based fencing to disk-based fencing using the installer” on page 444.
Migrating manually	See “To manually migrate from server-based fencing to disk-based fencing” on page 447.

To migrate from server-based fencing to disk-based fencing using the installer

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the SFCFSHA cluster is configured to use server-based fencing.

```
# vxfsadm -d
```

For example, if the SFCFSHA cluster uses server-based fencing, the output appears similar to the following:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
    * 0 (system01)
      1 (system02)
RFSM State Information:
    node 0 in state 8 (running)
    node 1 in state 8 (running)
```

- 3 On any node in the cluster, start the `installsfcsfsha` with the `-fencing` option.

```
# /opt/VRTS/install/installsfcsfsha -fencing
```

The `installsfcsfsha` starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 4 Confirm that you want to proceed with the I/O fencing configuration.
The installer verifies whether I/O fencing is configured in enabled mode.
- 5 Confirm that you want to reconfigure I/O fencing.
- 6 Review the I/O fencing configuration options that the program presents. Type **4** to migrate to disk-based I/O fencing.

Select the fencing mechanism to be configured in this
Application Cluster [1-4,q] **4**

- 7 From the list of coordination points that the installer presents, select the coordination points that you want to replace.

For example:

Select the coordination points you would like to remove
from the currently configured coordination points:

- 1) emc_clariion0_62
- 2) [10.209.80.197]:14250,[10.209.80.199]:14300
- 3) [10.209.80.198]:14250
- 4) All
- 5) None
- b) Back to previous menu

Enter the options separated by spaces: [1-5,b,q,?] (5)? **2 3**

- 8 Enter the total number of new coordination points.
- 9 Enter the total number of new coordinator disks.
- 10 From the list of available disks that the installer presents, select two disks which you want to configure as coordinator disks.

For example:

List of available disks:

- 1) emc_clariion0_61
- 2) emc_clariion0_65
- 3) emc_clariion0_66
- b) Back to previous menu

Select 2 disk(s) as coordination points. Enter the disk options
separated by spaces: [1-3,b,q]**2 3**

- 11 Verify and confirm the coordination points information for the fencing reconfiguration.

- 12** To migrate to disk-based fencing, select the I/O fencing mode as SCSI3.

```
Select the vxfen mode: [1-2,b,q,?] (1) 1
```

The installer initializes the coordinator disks and the coordinator disk group, and depots the disk group. Press **Enter** to continue.

- 13** Review the output as the installer prepares the `vxfenmode.test` file on all nodes and runs the `vxfenswap` script.

Note the location of the `vxfenswap.log` file which you can access in the event of any problem with the configuration process.

The installer cleans up the application cluster information from the CP servers.

- 14** If you want to send this installation information to Symantec, answer **y** at the prompt.

```
Would you like to send the information about this installation
to Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

- 15** After the migration is complete, verify the change in the fencing mode.

```
# vxfenadm -d
```

For example, after the migration from server-based fencing to disk-based fencing in the SFCFSA cluster:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 16** Verify the current coordination points that the vxfen driver uses.

```
# vxfenconfig -l
```

To manually migrate from server-based fencing to disk-based fencing

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the SFCFSHA cluster is online and uses server-based fencing.

```
# vxfenadm -d
```

For example, if SFCFSHA cluster uses server-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
  * 0 (system01)
  1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Make sure that you performed the following preparatory tasks to configure disk-based fencing:

- Identifying disks to use as coordinator disks
- Setting up coordinator disk group
- Creating I/O configuration files

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for the procedures.

- 4 Create a new `/etc/vxfenmode.test` file with the fencing configuration changes such as the scsi3 disk policy information.

Refer to the sample `vxfenmode` files in the `/etc/vxfen.d` folder.

- 5 From any node in the SFCFSHA cluster, start the `vxfenswap` utility:

```
# vxfenswap -g diskgroup [-n]
```

- 6 Review the message that the utility displays and confirm whether you want to commit the change.

- If you do not want to commit the new fencing configuration changes, press Enter or answer **n** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) n
```

The `vxfsnswap` utility rolls back the migration operation.

- If you want to commit the new fencing configuration changes, answer **y** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the
`/etc/vxfenmode.test` file to the `/etc/vxfenmode` file.

- 7 After the migration is complete, verify the change in the fencing mode.

```
# vxfsnadm -d
```

For example, after the migration from server-based fencing to disk-based fencing in the SFCFSHA cluster:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (system01)
    1 (system02)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 8 Verify the current coordination points that the `vxfsn` driver uses.

```
# vxfsnconfig -l
```

Migrating between fencing configurations using response files

Typically, you can use the response file that the installer generates after you migrate between I/O fencing configurations. Edit these response files to perform an automated fencing reconfiguration in the SFCFSHA cluster.

To configure I/O fencing using response files

- 1 Make sure that SFCFSHA is configured.
- 2 Make sure system-to-system communication is functioning properly.

- 3 Make sure that the SFCFSHA cluster is online and uses either disk-based or server-based fencing.

```
# vxvfenadm -d
```

For example, if SFCFSHA cluster uses disk-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
    * 0 (system01)
    1 (system02)
RFSM State Information:
    node 0 in state 8 (running)
    node 1 in state 8 (running)
```

For example, if the SFCFSHA cluster uses server-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
    * 0 (system01)
    1 (system02)
RFSM State Information:
    node 0 in state 8 (running)
    node 1 in state 8 (running)
```

- 4 Copy the response file to one of the cluster systems where you want to configure I/O fencing.

Review the sample files to reconfigure I/O fencing.

See [“Sample response file to migrate from disk-based to server-based fencing”](#) on page 450.

See [“Sample response file to migrate from server-based fencing to disk-based fencing”](#) on page 451.

See [“Sample response file to migrate from single CP server-based fencing to server-based fencing”](#) on page 451.

- 5 Edit the values of the response file variables as necessary.

See [“Response file variables to migrate between fencing configurations”](#) on page 451.

- 6 Start the I/O fencing reconfiguration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfcsfsha -responsefile /tmp/response_file
```

Where /tmp/response_file is the response file's full path name.

Sample response file to migrate from disk-based to server-based fencing

The following is a sample response file to migrate from disk-based fencing with three coordinator disks to server-based fencing with one CP server and two coordinator disks:

```
$CFG{disks_to_remove}=[ qw(emc_clariion0_62) ];  
$CFG{fencing_cps}=[ qw(10.198.89.251) ];  
$CFG{fencing_cps_ports}{"10.198.89.204"}=14250;  
$CFG{fencing_cps_ports}{"10.198.89.251"}=14250;  
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251 10.198.89.204) ];  
$CFG{fencing_ncp}=1;  
$CFG{fencing_option}=4;  
$CFG{opt}{configure}=1;  
$CFG{opt}{fencing}=1;  
$CFG{prod}="SFCFSHA60";  
$CFG{systems}=[ qw(system01 system02) ];  
$CFG{vcs_clusterid}=22462;  
$CFG{vcs_clustername}="cfs_cluster01";
```

Sample response file to migrate from server-based fencing to disk-based fencing

The following is a sample response file to migrate from server-based fencing with one CP server and two coordinator disks to disk-based fencing with three coordinator disks:

```
$CFG{fencing_disks}=[ qw(emc_clariion0_66) ];
$CFG{fencing_mode}="scsi3";
$CFG{fencing_ncp}=1;
$CFG{fencing_ndisks}=1;
$CFG{fencing_option}=4;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFCFS60";
$CFG{servers_to_remove}=[ qw([10.198.89.251]:14250) ];
$CFG{systems}=[ qw(system01 system02) ];
$CFG{vcs_clusterid}=42076;
$CFG{vcs_clustername}="cfs_cluster01";
```

Sample response file to migrate from single CP server-based fencing to server-based fencing

The following is a sample response file to migrate from single CP server-based fencing to server-based fencing with one CP server and two coordinator disks:

```
$CFG{fencing_disks}=[ qw(emc_clariion0_62 emc_clariion0_65) ];
$CFG{fencing_dgname}="fendg";
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{fencing_ncp}=2;
$CFG{fencing_ndisks}=2;
$CFG{fencing_option}=4;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFCFSHA60";
$CFG{systems}=[ qw(system01 system02) ];
$CFG{vcs_clusterid}=42076;
$CFG{vcs_clustername}="cfs_cluster01";
```

Response file variables to migrate between fencing configurations

[Table 15-7](#) lists the response file variables that specify the required information to migrate between fencing configurations for SFCFSHA.

Table 15-7 Response file variables specific to migrate between fencing configurations

Variable	List or Scalar	Description
CFG {fencing_option}	Scalar	Specifies the I/O fencing configuration mode. 1—Coordination Point Server-based I/O fencing 2—Coordinator disk-based I/O fencing 3—Disabled mode 4—Fencing migration when the cluster is online (Required)
CFG {fencing_reusedisk}	Scalar	If you migrate to disk-based fencing or to server-based fencing that uses coordinator disks, specifies whether to use free disks or disks that already belong to a disk group. 0—Use free disks as coordinator disks 1—Use disks that already belong to a disk group as coordinator disks (before configuring these as coordinator disks, installer removes the disks from the disk group that the disks belonged to.) (Required if your fencing configuration uses coordinator disks)
CFG {fencing_ncp}	Scalar	Specifies the number of new coordination points to be added. (Required)
CFG {fencing_ndisks}	Scalar	Specifies the number of disks in the coordination points to be added. (Required if your fencing configuration uses coordinator disks)

Table 15-7 Response file variables specific to migrate between fencing configurations (*continued*)

Variable	List or Scalar	Description
CFG {fencing_disks}	List	Specifies the disks in the coordination points to be added. (Required if your fencing configuration uses coordinator disks)
CFG {fencing_dgname}	Scalar	Specifies the disk group that the coordinator disks are in. (Required if your fencing configuration uses coordinator disks)
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the disk policy that the disks must use. (Required if your fencing configuration uses coordinator disks)
CFG {fencing_cps}	List	Specifies the CP servers in the coordination points to be added. (Required for server-based fencing)
CFG {fencing_cps_vips}{\$vip1}	List	Specifies the virtual IP addresses or the fully qualified host names of the new CP server. (Required for server-based fencing)
CFG {fencing_cps_ports}{\$vip}	Scalar	Specifies the port that the virtual IP of the new CP server must listen on. If you do not specify, the default value is 14250. (Optional)
CFG {servers_to_remove}	List	Specifies the CP servers in the coordination points to be removed.
CFG {disks_to_remove}	List	Specifies the disks in the coordination points to be removed

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy or group-based race policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

See [“About preferred fencing”](#) on page 58.

See [“How preferred fencing works”](#) on page 59.

To enable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfsadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To enable system-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as System.

```
# haclus -modify PreferredFencingPolicy System
```

- Set the value of the system-level attribute FencingWeight for each node in the cluster.

For example, in a two-node cluster, where you want to assign system01 five times more weight compared to system02, run the following commands:

```
# hasys -modify system01 FencingWeight 50  
# hasys -modify system02 FencingWeight 10
```

- Save the VCS configuration.

```
# haconf -dump -makero
```

- 4 To enable group-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

```
# haclus -modify PreferredFencingPolicy Group
```

- Set the value of the group-level attribute Priority for each service group. For example, run the following command:

```
# hagrps -modify service_group Priority 1
```

Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the VCS configuration.

```
# haconf -dump -makero
```

- 5 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute PreferredFencingPolicy as Disabled.

```
# haconf -makerw
```

```
# haclus -modify PreferredFencingPolicy Disabled
```

```
# haconf -dump -makero
```

About VXFEN tunable parameters

The section describes the VXFEN tunable parameters and how to reconfigure the VXFEN module.

[Table 15-8](#) describes the tunable parameters for the VXFEN driver.

Table 15-8 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_debug_sz	Size of debug log in bytes - Values - Default: 131072 (128 KB) - Minimum: 65536 (64 KB) - Maximum: 524288 (512 KB)
vxfen_max_delay	Specifies the maximum number of seconds that the smaller sub-cluster waits before racing with larger sub-clusters for control of the coordinator disks when a network partition occurs. This value must be greater than the vxfen_min_delay value. - Values - Default: 60 - Minimum: 1 - Maximum: 600
vxfen_min_delay	Specifies the minimum number of seconds that the smaller sub-cluster waits before racing with larger sub-clusters for control of the coordinator disks when a network partition occurs. This value must be smaller than or equal to the vxfen_max_delay value. - Values - Default: 1 - Minimum: 1 - Maximum: 600
vxfen_vxfnd_tmt	Specifies the time in seconds that the I/O fencing driver VxFEN waits for the I/O fencing daemon VXFEND to return after completing a given task. - Values - Default: 60 - Minimum: 10 - Maximum: 600

Table 15-8 VXFEN tunable parameters (*continued*)

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
panic_timeout_offst	Specifies the time in seconds based on which the I/O fencing driver VxFEN computes the delay to pass to the GAB module to wait until fencing completes its arbitration before GAB implements its decision in the event of a split-brain. You can set this parameter in the <code>vxfenmode</code> file and use the <code>vxfenadm</code> command to check the value. Depending on the <code>vxfen_mode</code>, the GAB delay is calculated as follows: ■ For <code>scsi3</code> mode: $1000 * (\text{panic\_timeout\_offst} + \text{vxfen\_max\_delay})$ ■ For customized mode: $1000 * (\text{panic\_timeout\_offst} + \max(\text{vxfen\_vxwnd\_tmt}, \text{vxfen\_loser\_exit\_delay}))$ ■ Default: 10

In the event of a network partition, the smaller sub-cluster delays before racing for the coordinator disks. The time delay allows a larger sub-cluster to win the race for the coordinator disks. The `vxfen_max_delay` and `vxfen_min_delay` parameters define the delay in seconds.

Configuring the VXFEN module parameters

After adjusting the tunable kernel driver parameters, you must reconfigure the VXFEN module for the parameter changes to take effect.

The following example procedure changes the value of the `vxfen_min_delay` parameter.

On each Solaris node, edit the file `/kernel/drv/vxfen.conf` to change the value of the vxfen driver tunable global parameters, `vxfen_max_delay` and `vxfen_min_delay`.

Note: You must restart the VXFEN module to put any parameter change into effect.

To configure the VxFEN parameters and reconfigure the VxFEN module

- 1 Edit the file `/kernel/drv/vxfen.conf` to change the `vxfen_min_delay` value to 30.

The following VXFEN example displays the content of the default file `/kernel/drv/vxfen.conf` before changing the `vxfen_min_delay` parameter:

```
#
# VXFEN configuration file
#
name="vxfen" parent="pseudo" instance=0 dbg_log_size=65536
vxfen_max_delay=60 vxfen_min_delay=1;
```

After editing the file to change the `vxfen_min_delay` value to 30, the default file `/kernel/drv/vxfen.conf` contains the following values:.

```
#
# VXFEN configuration file
#
name="vxfen" parent="pseudo" instance=0 dbg_log_size=65536
vxfen_max_delay=60 vxfen_min_delay=30;
```

After reviewing the edits that you made to the default file, close and save the file.

- 2 Stop VCS.

```
# hastop -local
```

- 3 Unconfigure the VXFEN module:

```
# vxfenconfig -U
```

- 4 Determine the VXFEN module ID:

```
# /usr/sbin/modinfo | grep -i vxfen
```

The module ID is the number in the first column of the output.

- 5 Unload the VXFEN module, using the module ID you determined:

```
# /usr/sbin/modunload -i module_ID
```

- 6 Run the `update_drv` command to re-read the `/kernel/drv/vxfen.conf` file.

```
# /usr/sbin/update_drv vxfen
```

Note:

The `modunload` command has often been used on driver modules to force the system to reread the associated driver configuration file. While this procedure and command works in Solaris 9, this behavior may fail in later releases. The supported method for rereading the driver configuration file for systems running Solaris 10 or later is through the `update_drv` command. For additional information, refer to `update_drv(1M)`.

- 7 Configure the VXFEN module:

```
# vxfenconfig -c
```

- 8 Start VCS.

```
# hstart
```

Administering SFCFSHA global clusters

This section provides instructions for the following global cluster administration tasks:

- About setting up a fire drill
See [“About setting up a disaster recovery fire drill”](#) on page 459.
- Configuring the fire drill service group using the wizard
See [“About configuring the fire drill service group using the Fire Drill Setup wizard”](#) on page 460.
- Verifying a successful fire drill
See [“Verifying a successful fire drill”](#) on page 462.
- Scheduling a fire drill
See [“Scheduling a fire drill”](#) on page 463.

About setting up a disaster recovery fire drill

The disaster recovery fire drill procedure tests the fault-readiness of a configuration by mimicking a failover from the primary site to the secondary site. This procedure is done without stopping the application at the primary site and disrupting user

access, interrupting the flow of replicated data, or causing the secondary site to need resynchronization.

The initial steps to create a fire drill service group on the secondary site that closely follows the configuration of the original application service group and contains a point-in-time copy of the production data in the Replicated Volume Group (RVG). Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over and come online at the secondary site, should the need arise. Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online.

You must conduct a fire drill only at the secondary site; do not bring the fire drill service group online on the node hosting the original application.

Before you perform a fire drill in a disaster recovery setup that uses VVR, perform the following steps:

- Set the value of the ReuseMntPt attribute to 1 for all Mount resources.
- Configure the fire drill service group.
See [“About configuring the fire drill service group using the Fire Drill Setup wizard”](#) on page 460.
- After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 0 for all Mount resources.

VCS also supports HA fire drills to verify a resource can fail over to another node in the cluster.

Note: You can conduct fire drills only on regular VxVM volumes; volume sets (vset) are not supported.

VCS provides hardware replication agents for array-based solutions, such as Hitachi Truecopy, EMC SRDF, and so on . If you are using hardware replication agents to monitor the replicated data clusters, refer to the VCS replication agent documentation for details on setting up and configuring fire drill.

About configuring the fire drill service group using the Fire Drill Setup wizard

Use the Fire Drill Setup Wizard to set up the fire drill configuration.

The wizard performs the following specific tasks:

- Creates a Cache object to store changed blocks during the fire drill, which minimizes disk space and disk spindles required to perform the fire drill.

- Configures a VCS service group that resembles the real application group.

The wizard works only with application groups that contain one disk group. The wizard sets up the first RVG in an application. If the application has more than one RVG, you must create space-optimized snapshots and configure VCS manually, using the first RVG as reference.

You can schedule the fire drill for the service group using the `fdsched` script.

See [“Scheduling a fire drill”](#) on page 463.

Running the fire drill setup wizard

To run the wizard

- 1 Start the RVG Secondary Fire Drill wizard on the VVR secondary site, where the application service group is offline and the replication group is online as a secondary:

```
# /opt/VRTSvcs/bin/fdsetup
```

- 2 Read the information on the Welcome screen and press the **Enter** key.
- 3 The wizard identifies the global service groups. Enter the name of the service group for the fire drill.
- 4 Review the list of volumes in disk group that could be used for a space-optimized snapshot. Enter the volumes to be selected for the snapshot. Typically, all volumes used by the application, whether replicated or not, should be prepared, otherwise a snapshot might not succeed.

Press the **Enter** key when prompted.

- 5 Enter the cache size to store writes when the snapshot exists. The size of the cache must be large enough to store the expected number of changed blocks during the fire drill. However, the cache is configured to grow automatically if it fills up. Enter disks on which to create the cache.

Press the **Enter** key when prompted.

- 6 The wizard starts running commands to create the fire drill setup.

Press the **Enter** key when prompted.

The wizard creates the application group with its associated resources. It also creates a fire drill group with resources for the application (Oracle, for example), the CFSSMount, and the RVGSnapshot types.

The application resources in both service groups define the same application, the same database in this example. The wizard sets the FireDrill attribute for the application resource to 1 to prevent the agent from reporting a concurrency violation when the actual application instance and the fire drill service group are online at the same time.

About configuring local attributes in the fire drill service group

The fire drill setup wizard does not recognize localized attribute values for resources. If the application service group has resources with local (per-system) attribute values, you must manually set these attributes after running the wizard.

Verifying a successful fire drill

Bring the fire drill service group online on a node that does not have the application running. Verify that the fire drill service group comes online. This action validates that your disaster recovery solution is configured correctly and the production service group will fail over to the secondary site in the event of an actual failure (disaster) at the primary site.

If the fire drill service group does not come online, review the VCS engine log to troubleshoot the issues so that corrective action can be taken as necessary in the production service group.

You can also view the fire drill log, located at `/tmp/fd-servicegroup.pid`

Remember to take the fire drill offline once its functioning has been validated. Failing to take the fire drill offline could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Scheduling a fire drill

You can schedule the fire drill for the service group using the `fdsched` script. The `fdsched` script is designed to run only on the lowest numbered node that is currently running in the cluster. The scheduler runs the command `hagrp - online firedrill_group -any` at periodic intervals.

To schedule a fire drill

- 1 Add the file `/opt/VRTSvcs/bin/fdsched` to your crontab.
- 2 To make fire drills highly available, add the `fdsched` file to each node in the cluster.

Administering sites and remote mirrors


This chapter includes the following topics:

- [About sites and remote mirrors](#)
- [Making an existing disk group site consistent](#)
- [Configuring a new disk group as a Remote Mirror configuration](#)
- [Fire drill — testing the configuration](#)
- [Changing the site name](#)
- [Administering the Remote Mirror configuration](#)
- [Examples of storage allocation by specifying sites](#)
- [Displaying site information](#)
- [Failure and recovery scenarios](#)

About sites and remote mirrors

In a Remote Mirror configuration (also known as a campus cluster or stretch cluster) the hosts and storage of a cluster that would usually be located in one place, are instead divided between two or more sites. These sites are typically connected via a redundant high-capacity network that provides access to storage and private link communication between the cluster nodes.

[Figure 16-1](#) shows a typical two-site remote mirror configuration.

Figure 16-1 Example of a two-site remote mirror configuration


If a disk group is configured across the storage at the sites, and inter-site communication is disrupted, there is a possibility of a serial split brain condition arising if each site continues to update the local disk group configuration copies.

VxVM provides mechanisms for dealing with the serial split brain condition, monitoring the health of a remote mirror, and testing the robustness of the cluster against various types of failure (also known as fire drill).

For applications and services to function correctly at a site when other sites have become inaccessible, at least one complete plex of each volume must be configured at each site (site-based allocation), and the consistency of the data in the plexes at each site must be ensured (site consistency).

By tagging disks with site names, storage can be allocated from the correct location when creating, resizing or relocating a volume, and when changing a volume's layout.

[Figure 16-2](#) shows an example of a site-consistent volume with two plexes configured at each of two sites.

Figure 16-2 Site-consistent volume with two plexes at each of two sites

The storage for plexes P1 and P2 is allocated storage that is tagged as belonging to site A, and the storage for plexes P3 and P4 is allocated storage that is tagged as belonging to site B.


Although not shown in this figure, DCO log volumes are also mirrored across the sites, and disk group configuration copies are distributed across the sites.

Site consistency means that the data in the plexes for a volume must be consistent at each site. The site consistency of a volume is ensured by detaching a site when its last complete plex fails at that site. If a site fails, all its plexes are detached and the site is said to be detached. If site consistency is not on, only the plex that fails is detached. The remaining volumes and their plexes on that site are not detached.

To enhance read performance, VxVM will service reads from the plexes at the local site where an application is running if the `siteread` read policy is set on a volume. Writes are written to plexes at all sites.

[Figure 16-3](#) shows a configuration with remote storage only that is also supported.

Figure 16-3 Example of a two-site configuration with remote storage only



About site-based allocation

Site-based allocation policies are enforced by default in a site-configured disk group. Site-based allocation requires that each volume has at least one plex at each site that is configured in the disk group. When a new volume is created in a site-configured disk group, the `allsites` attribute is set to `on` by default. The `allsites` attribute indicates that the volume must have at least one plex on each configured site in the disk group. For new volumes, the read policy is set to `siteread` by default.

If mirroring across sites is not required, or is not possible (as is the case for RAID-5 volumes), specify the `allsites=off` attribute to the `vxassist` command. If sites are configured in the disk group, a plex will always be confined to a site and will not span across sites. This enforcement cannot be overridden.

Before adding a new site to a disk group, be sure to meet the following requirements:

- Disks from the site being added (site tagged) are present or added to the disk group.
- Each existing volume with `allsites` set in the disk group must have at least one plex at the site being added. If this condition is not met, the command to

add the site to the disk group fails. If the `-f` option is specified, the command does not fail, but instead it sets the `allsites` attribute for the volume to `off`.

About site consistency

Site consistency means that at any point in time, the data at each site is consistent with the application for a given set of volumes. A site-consistent volume must have at least one plex, or mirror, on each configured site in the disk group. The site consistency is ensured by detaching a site when a site-consistent volume loses its last complete plex on that site. The `site detach` detaches all the plaxes on that site and also disallows further configuration updates to the configuration copies on that site. Turn on this behavior by setting the `siteconsistent` attribute to `on` on the desired volumes.

If you set the `siteconsistent` attribute to `off`, only the plex that fails is detached. The plaxes for the remaining volumes on that site are not detached.

The `siteconsistent` attribute is also present at the disk group level and can be used to turn on or off the site consistency functionality in the disk group boundary. In addition, if you turn on the `siteconsistent` attribute for a disk group, each new volume created in the disk group inherits the site consistency of the disk group, by default. Setting the `siteconsistent` attribute on a disk group does not affect `siteconsistent` attributes for existing volumes. You can also control the site consistency on individual volumes.

By default, a volume inherits the value that is set on its disk group.

By default, creating a site-consistent volume also creates an associated version 20 DCO volume, and enables Persistent FastResync on the volume. This allows faster recovery of the volume during the reattachment of a site.

See [“Configuring site consistency on a volume”](#) on page 476.

Before setting site consistency on a disk group, be sure to meet the following requirements:

- A license enabling the Site Awareness feature must be installed on all the hosts in the Remote Mirror configuration.
- At least two sites must be configured in the disk group before site consistency is turned on.
See [“Making an existing disk group site consistent”](#) on page 470.
- All the disks in a disk group must be registered to one of the sites before you can set the `siteconsistent` attribute on the disk group.

About site tags

In a Remote Mirror configuration, each storage device in the disk group must be tagged with site information. The site tag indicates to which site the device is associated. VxVM provides a facility to tag VxVM-initialized disks with an arbitrary name-value pair. The tag name `site` is reserved by VxVM and is used to identify the site information of tagged disks. The command `vxdisk settag` can be used to tag multiple disks or all disks from an enclosure or disks from multiple enclosures. The tagging can be restricted to disks in a disk group by specifying the disk group with the command.

You can use automatic site tagging to assign site tags to disks when adding them to a disk group. When automatic site tagging is on, newly added disks or LUNs inherit the site tag from the site-enclosure mapping stored in the disk group. To use automatic site tagging, turn on automatic site tagging for a disk group, and then assign the site names to the enclosures in the disk group. Any disks or LUNs in that disk group inherit the tag from the enclosure to which they belong.

About the site read policy

To enhance read performance, VxVM will service reads from the plexes at the local site where an application is running, if the `siteread` read policy is set on a volume. Writes are written to plexes at all sites. By tagging hosts with site information, VxVM identifies which hosts belong to which site. Reads initiated by a host from one site are then satisfied by disks which are tagged with the same site. Tagging hosts and disks with correct site information gives you maximum read performance when `siteread` read policy is used.

If a license enabling the Site Awareness feature is installed on all the hosts in the Remote Mirror configuration, the disk group is configured for site consistency with several sites enabled, and the `allsites=on` attribute is specified for a volume, the default read policy is `siteread`.

If the `siteread` policy is not set, use the following command to set the read policy to `siteread` policy on a volume:

```
# vxvol [-g diskgroup] rdpol siteread volume
```

This command has no effect if a site name has not been set for the host.

See [“Changing the read policy for mirrored volumes”](#) on page 209.

Making an existing disk group site consistent

The site consistency feature requires that a license enabling the site awareness feature has been installed on all hosts at all sites that participate in the configuration.

To make an existing disk group site consistent

- 1 Ensure that the disk group is updated to at least version 140, by running the `vxchg upgrade` command on it:

```
# vxchg upgrade diskgroup
```

- 2 On each host that can access the disk group, define the site name:

```
# vxctl set site=sitename
```

- 3 Tag all the disks in the disk group with the appropriate site name:

```
# vxdisk [-g diskgroup] settag site=sitename  
disk1  
disk2
```

Or, to tag all the disks in a specified enclosure, use the following command:

```
# vxdisk [-g diskgroup] settag site=sitename  
encl:encl_name
```

- 4 Use the `vxchg move` command to move any unsupported RAID-5 volumes to another disk group. Alternatively, use the `vxassist convert` commands to convert the volumes to a supported layout such as `mirror` or `mirror-stripe`. You can use the `site` and `mirror=site` storage allocation attribute to ensure that the plexes are created on the correct storage.
- 5 Use the `vxevac` command to ensure that the volumes have at least one plex at each site. You can use the `site` and `mirror=site` storage allocation attribute to ensure that the plexes are created on the correct storage.
- 6 Register a site record for each site with the disk group:

```
# vxchg -g diskgroup addsite sitename
```

- 7 Turn on site consistency for the disk group:

```
# vxchg -g diskgroup set siteconsistent=on
```

- 8 Turn on the `allsites` flag for the volume which requires data replication to each site:

```
# vxvol [-g diskgroup] set allsites=on volume
```

- 9 Turn on site consistency for each existing volume in the disk group for which siteconsistency is needed. You also need to attach `DCOV20` if it is not attached already. `DCOV20` is required to ensure that site detach and reattach are instantaneous.

```
# vxvol [-g diskgroup] set siteconsistent=on volume ...
```

Configuring a new disk group as a Remote Mirror configuration

Note: The Remote Mirror feature requires that a license enabling the Site Awareness feature has been installed on all hosts at all sites that participate in the configuration.

This section describes setting up a new disk group. To configure an existing disk group as a Remote Mirror configuration, additional steps may be required.

See [“Making an existing disk group site consistent”](#) on page 470.

Setting up a new disk group for a Remote Mirror configuration

- 1 Define the site name for each host that can access the disk group.

```
# vxdctl set site=sitename
```

To verify the site name assigned to the host, use the following command:

```
# vxdctl list
```

- 2 Create the disk group with storage from each site.
- 3 Register a site record to the disk group, for each site.

```
# vxdg -g diskgroup [-f] addsite sitename
```

- 4 Do one of the following:
 - To tag all disks regardless of the disk group, do the following:
Assign a site name to the disks or enclosures. You can set site tags at the disk level, or at the enclosure level. If you specify one or more enclosures,

the site tag applies to the disks in that enclosure that are within the disk group. Enter the following command:

```
# vxdisk [-g diskgroup] settag site=sitename \  
disk disk1... |encl:encl_name encl:encl_name1...
```

where the disks can be specified either by the disk access name or the disk media name.

- To autotag new disks added to the disk group based on the enclosure to which they belong, perform the following steps in the order presented. These steps are limited to disks in a single group.
 - Set the `autotagging` policy to `on` for the disk group, if required. Automatic tagging is the default setting, so this step is only required if the autotagging policy was previously disabled. To turn on autotagging, enter the following command:

```
# vxdg [-g diskgroup] set autotagging=on
```

- Add site-enclosure mapping information to the diskgroup for each site-enclosure combination. Enter the following command:

```
# vxdg [-g diskgroup] settag encl:encl_name1 site=sitename1
```

As a result of this command, all disks of enclosure `encl_name1` in the specified disk group are tagged with site information.

5 Turn on the site consistency requirement for a disk group:

```
# vxdg -g diskgroup set siteconsistent=on
```

Fire drill — testing the configuration

Warning: To avoid potential loss of service or data, it is recommended that you do not use these procedures on a live system.

After validating the consistency of the volumes and disk groups at your sites, you should validate the procedures that you will use in the event of the various possible types of failure. A fire drill lets you test that a site can be brought up cleanly during recovery from a disaster scenario such as site failure.

Simulating site failure

To simulate the failure of a site, use the following command to detach all the devices at a specified site:

```
# vxdg -g diskgroup [-f] detachsite sitename
```

The `-f` option must be specified if any plexes configured on storage at the site are currently online.

After the site is detached, the application should run correctly on the available site. This step verifies that the primary site is fine. Continue the fire drill by verifying the secondary site.

Verifying the secondary site

After detaching the site from primary site, verify whether the application starts correctly on a secondary site. The fire drill ensures that the application can run on the secondary if disaster strikes the primary site. These procedures assume that the application is running correctly before the fire drill operation begins.

To verify the secondary site, import the detached site on a different host using the following command:

```
# vxdg -o site=sitename import dgname
```

Then start the application. If the application runs correctly on the secondary site, this step verifies the integrity of the secondary site.

Recovery from simulated site failure

After verifying the data on the secondary for a simulated site failure, deport the disk group from the secondary site. Then reattach the site back to the primary host.

Use the following commands to reattach a site and recover the disk group:

```
# vxdg -g diskgroup [-o overridesb] reattachsite sitename  
# vxrecover -g diskgroup
```

It may be necessary to specify the `-o overridesb` option if a serial split-brain condition is indicated.

Changing the site name

You can change the site name, or tag, that is used to identify each site in a Remote Mirror configuration. Renaming the site changes the site record in the disk group.

The site name is also changed for all of the disks and enclosures that are tagged with the existing site name.

After you rename a site, you need to explicitly change the site name for each host that belongs to that site.

See [“Resetting the site name for a host”](#) on page 474.

To rename the site

- ◆ Specify the new site name as follows:

```
# vxdg [-g diskgroup] renamesite old_sitename new_sitename
```

Resetting the site name for a host

If you rename a site, you need to explicitly set each host to refer to the new site name.

To reset a site name for a host

- 1 Remove the site name from a host:

```
# vxdctl [-F] unset site
```

The `-F` option is required if any imported disk groups are registered to the site.

- 2 Set the new site name for the host.

```
# vxdctl set site=sitename
```

The name that has been assigned to a site is stored in the `/etc/vx/volboot` file.

Administering the Remote Mirror configuration

After the Remote Mirror site is configured, refer to the following sections for additional tasks to maintain the configuration.

Configuring site tagging for disks or enclosures

To set up a Remote Mirror configuration, specify to which site each storage device in the disk group belongs. Assign a site tag to one or more disks or enclosures. If the disk or enclosure does not belong to a disk group, you must use this method to assign a site tag.

To tag disks or enclosures with a site name

- ◆ Assign a site name to one or more disks or enclosures, using the following command:

```
# vxdisk [-g diskgroup] settag site=sitename \  
disk disk1...|encl:encl_name encl:encl_name1...
```

where the disks can be specified either by the disk access name or the disk media name.

To display the disks or enclosures registered to a site

- ◆ To check which disks or enclosures are registered to a site, use the following command:

```
# vxdisk [-g diskgroup] listtag
```

To remove the site tag from a disk or enclosure

- ◆ To remove the site tag from a disk or enclosure, use the following command:

```
# vxdisk rmtag site=sitename \  
disk disk1...|encl:encl_name encl:encl_name1...
```

Configuring automatic site tagging for a disk group

Configure automatic site tagging if you want disks or LUNs to inherit the tag from the enclosure. After you turn on automatic site tagging for a disk group, assign the site names to the enclosures in the disk group. Any disks or LUNs added to that disk group inherit the tag from the enclosure to which they belong.

To configure automatic site tagging for a disk group

- 1 Set the autotagging policy to **on** for the disk group. Automatic tagging is the default setting, so this step is only required if the autotagging policy was previously disabled.

To turn on autotagging, use the following command:

```
# vxdg [-g diskgroup] set autotagging=on
```

- 2 Assign the site name to an enclosure within the disk group, using the following command:

```
# vxdg [-g diskgroup] settag encl:encl_name site=sitename
```

To list the site tags for a disk group

- ◆ To list the site tags for a disk group, use the following command:

```
# vxdg [-g diskgroup] listtag site=sitename
```

To remove a site tag from an enclosure or a disk group

- ◆ To remove a site tag from a disk group, use the following command:

```
# vxdg [-g diskgroup] rmtag [encl:encl_name] site=sitename
```

Configuring site consistency on a volume

To set the site consistency requirement when creating a volume, specify the `siteconsistent` attribute to the `vxassist make` command, for example:

```
# vxassist [-g diskgroup] make volume size \
  nmirror=4 siteconsistent={on|off}
```

By default, a volume inherits the value that is set on its disk group.

By default, creating a site-consistent volume also creates an associated version 20 DCO volume, and enables Persistent FastResync on the volume. This allows faster recovery of the volume during the reattachment of a site.

To turn on the site consistency requirement for an existing volume, use the following form of the `vxvol` command:

```
# vxvol [-g diskgroup] set siteconsistent=on volume
```

To turn off the site consistency requirement for a volume, use the following command:

```
# vxvol [-g diskgroup] set siteconsistent=off volume
```

The `siteconsistent` attribute and the `allsites` attribute must be set to `off` for RAID-5 volumes in a site-consistent disk group.

Examples of storage allocation by specifying sites

Table 16-1 shows examples of how to use sites with the `vxassist` command to allocate storage. These examples assume that the disk group, `ccdg`, has been enabled for site consistency with disks configured at two sites, `site1` and `site2`. Also, `ccdg01`, `ccdg02`, and `ccdg03` are dm names of disks tagged with site `site1`. `ccdg09`, `ccdg10`, and `ccdg11` are dm names of disks tagged with site `site2`.

Table 16-1 Examples of storage allocation by specifying sites

Command	Description
# vxassist -g ccdg make vol 2g \ nmirror=2	Create a volume with one mirror at each site. The <code>nmirror</code> keyword is optional. If the <code>nmirror</code> keyword is specified, it must equal the number of sites.
# vxassist -g ccdg -o ordered \ make vol 2g \ layout=mirror-stripe ncol=3 \ ccdg01 ccdg02 ccdg03 ccdg09 \ ccdg10 ccdg11	Create a mirrored-stripe volume specifying allocation order to validate redundancy across the sites. The named disks must be tagged with the appropriate site name, and there must be sufficient disks at each site to create the volume.
# vxassist -g ccdg make vol 2g \ nmirror=2 ccdg01 ccdg09	Create a volume with one mirror on each of the named disks. The named disks must be tagged with the appropriate site name, and there must be sufficient disks at each site to create the volume.
# vxassist -g ccdg make vol 2g \ nmirror=2 siteconsistent=off \ allsites=off	Create a mirrored volume that is not site consistent. Both mirrors can be allocated from any available storage in the disk group, but the storage for each mirror is confined to a single site.
# vxassist -g ccdg make vol 2g \ nmirror=2 site:site2 \ siteconsistent=off \ allsites=off	Create a mirrored volume that is not site consistent. Both mirrors are allocated from any available storage in the disk group that is tagged as belonging to <code>site2</code> .
# vxassist -g ccdg make vol 2g \ nmirror=2 \!site:site1 \ siteconsistent=off \ allsites=off	Create a mirrored volume that is not site consistent. Both mirrors are allocated from any available storage in the disk group that is tagged as not belonging to <code>site1</code> . Note: The <code>!</code> character is a special character in some shells. This example shows how to escape it in a bash shell.

Table 16-1 Examples of storage allocation by specifying sites (*continued*)

Command	Description
<pre># vxassist -g ccdg mirror vol \ site:sitel</pre>	Add a mirror at a specified site. The command fails if there is insufficient storage available at the site. This command does not affect the <code>allsites</code> or <code>siteconsistent</code> of a volume.
<pre># vxassist -g ccdg remove \ mirror vol site:sitel</pre>	Remove a mirror from a volume at a specified site. If the volume has the <code>allsites</code> attribute set to on, the command fails if this would remove the last remaining plex at a site.
<pre># vxassist -g ccdg growto vol \ 4g</pre>	Grow a volume. Each mirror of a volume is grown using the same site storage to which it belongs. If there is not enough storage to grow a mirror on each site, the command fails.

Displaying site information

To display the site name for a host

- ◆ To determine to which site a host belongs, use the following command on the host:

```
# vxdctl list | grep siteid
siteid: building1
```

To display the disks or enclosures registered to a site

- ◆ To check which disks or enclosures are registered to a site, use the following command:

```
# vxdisk [-g diskgroup] listtag
```

To display the setting for automatic site tagging for a disk group

- ◆ To determine whether automatic site tagging is on for a disk group, use the following command:

```
# vxprint -g diskgroup -F"%autotagging" diskgroup
```

To verify whether site consistency has been enabled for a disk group

- ◆ To verify whether site consistency has been enabled for a disk group, use the following command:

```
# vxdg list diskgroup | grep siteconsistent
flags: siteconsistent
```

To verify whether site consistency has been enabled for a volume

- ◆ To verify whether site consistency has been enabled for a volume, use the following command:

```
# vxprint -g diskgroup -F"%siteconsistent" vol
```

To identify which site a plex or mirror is allocated from

- ◆ To identify which site a plex or mirror is allocated from, use the following command:

```
# vxprint -g diskgroup -F"%site" plex
```

To list the site tags for a disk group

- ◆ To list the site tags for a disk group, use the following command:

```
# vxdg [-g diskgroup] listtag site=sitename
```

Failure and recovery scenarios

[Table 16-2](#) lists the possible failure scenarios and recovery procedures for the Remote Mirror feature.

Table 16-2 Failure scenarios and recovery procedures

Failure scenario	Recovery procedure
Disruption of network link between sites.	See “Recovering from a loss of site connectivity” on page 480.
Failure of hosts at a site.	See “Recovering from host failure” on page 481.
Failure of storage at a site.	See “Recovering from storage failure” on page 481.
Failure of both hosts and storage at a site.	See “Recovering from site failure” on page 481.

Table 16-2 Failure scenarios and recovery procedures (*continued*)

Failure scenario	Recovery procedure
Disruption of connectivity to storage at each others' remote sites for hosts on all sites.	See "Recovering from disruption of connectivity to storage at the remote sites from hosts on all sites" on page 482.
Disruption of connectivity to storage at all sites for hosts at a site	See "Recovering from disruption to connectivity to storage at all sites from the hosts at a site" on page 482.

Recovering from a loss of site connectivity

Warning: To avoid a potential loss of data, it is recommended that you configure Veritas Cluster Server to handle network split-brain.

If the network links between the sites are disrupted, the application environments may continue to run in parallel, and this may lead to inconsistencies between the disk group configuration copies at the sites. If the parallel instances of an application issue writes to volumes, an unrecoverable data loss may occur and manual intervention is needed. To avoid data loss, it is recommended that you configure the VCS fencing mechanism to handle network split-brain situations.

If VCS fencing is not used, serial split-brain condition may occur. When connectivity between the sites is restored, a serial split-brain condition will be detected between the sites. One site must be chosen as having the preferred version of the data and the disk group configuration copies. The data from the chosen site is resynchronized to other the site. If new writes are issued to volumes after the network split, they are overwritten with the data from the chosen site. The configuration copies at the other sites are updated from the copies at the chosen site.

At the chosen site, use the following commands to reattach a site and recover the disk group:

```
# vxdbg -g diskgroup -o overridessb reattachsite sitename
# vxrecover -g diskgroup
```

In the case that the host systems are configured at a single site with only storage at the remote sites, the usual resynchronization mechanism of VxVM is used to recover the remote plexes when the storage comes back on line.

Recovering from host failure

If one or more cluster nodes fail at a site, but the storage remains online, this is handled either by VCS failover in the case of the Storage Foundation HA product, or by node takeover in the case that the node was the master for a shared disk group as supported by the Storage Foundation Cluster File System software.

Recovering from storage failure

If storage fails at a site, the plexes that are configured on that storage are detached locally if a site-consistent volume still has other mirrors available at the site. The hot-relocation feature of VxVM will attempt to recreate the failed plexes on other available storage in the disk group. If no plexes of a site-consistent volume remain in operation at a site, and hot-relocation cannot recreate the plexes at that site, the site is detached. Because site connectivity has not been lost, applications running on hosts at the site can still access data at the other sites.

When the storage comes back online, the `vxattachd` reattaches the site automatically.

See [“Automatic site reattachment”](#) on page 482.

If the `vxattachd` is not running, use the following commands to reattach a site and recover the disk group:

```
# vxdg -g diskgroup reattachsite sitename
# vxrecover -g diskgroup
```

For more information about recovering a disk group, refer to the *Veritas Storage Foundation and High Availability Troubleshooting Guide*.

Recovering from site failure

If all the hosts and storage fail at a site, use the following commands to reattach the site after it comes back online, and to recover the disk group:

```
# vxdg -g diskgroup [-o overridesb] reattachsite sitename
# vxrecover -g diskgroup
```

The `-o overridesb` option is only required if a serial split-brain condition is indicated. A serial split-brain condition may happen if the site was brought back up while the private network link was inoperative. This option updates the configuration database on the reattached site with the consistent copies at the other sites.

For more information about recovering a disk group, refer to the *Veritas Storage Foundation and High Availability Troubleshooting Guide*.

Recovering from disruption of connectivity to storage at the remote sites from hosts on all sites

In this scenario, hosts at the sites lose connectivity to the storage at each others' sites. For example, hosts in Site A lose connectivity to the storage at Site B, and hosts in site B lose connectivity to the storage at Site A. For example, this might occur because the link between the fibre channel (FC) switches went down.

In this case, one of the sites is detached. Hosts in the detached site see I/O failures. Applications on the hosts in the detached site should be failed over to hosts on the other site.

To recover, reconnect the FC links. If the `vxattachd` daemon is running, it automatically reattaches the detached site and initiates recovery. Otherwise, manually run the `vxreattach` and `vxrecover` commands to bring the site back to ACTIVE.

See [“Recovering from storage failure”](#) on page 481.

Recovering from disruption to connectivity to storage at all sites from the hosts at a site

In this scenario, hosts at a site lose connectivity to the storage at all sites. For example, hosts in Site A lose connectivity to the storage at both Site A and Site B.

In this case, no site is detached. I/O fails on the hosts that lose connectivity to the storage at all sites. The application should be failed over to hosts on the other site.

Restore the connectivity to the storage, and then bring the applications online.

Automatic site reattachment

The automatic site reattachment daemon, `vxattachd`, provides automatic reattachment of sites. The `vxattachd` daemon uses the `vxnotify` mechanism to monitor storage coming back online on a site after a previous failure, and to restore redundancy of mirrors across sites.

If the hot-relocation daemon, `vxrelocd`, is running, `vxattachd` attempts to reattach the site, and allows `vxrelocd` to try to use the available disks in the disk group to relocate the failed subdisks. If `vxrelocd` succeeds in relocating the failed subdisks, it starts the recovery of the plexes at the site. When all the plexes have been recovered, the plexes are put into the ACTIVE state, and the state of the site is set to ACTIVE.

If `vxrelocd` is not running, `vxattachd` reattaches a site only when all the disks at that site become accessible. After reattachment succeeds, `vxattachd` sets the site

state to ACTIVE, and initiates recovery of the plexes. When all the plexes have been recovered, the plexes are put into the ACTIVE state.

Note: `vxattachd` does not try to reattach a site that you have explicitly detached by using the `vxdbg detachsite` command.

The automatic site reattachment feature is enabled by default. The `vxattachd` daemon uses email to notify `root` of any attempts to reattach sites and to initiate recovery of plexes at those sites.

To send mail to other users, add the user name to the line that starts `vxattachd` in the `/lib/svc/method/vxvm-recover` startup script and run the `svcadm refresh vxvm/vxvm-recover` command (for Solaris 10 onward), or `/etc/init.d/vxvm-recover` and reboot the system (for OS releases before Solaris 10).

If you do not want a site to be recovered automatically, kill the `vxattachd` daemon, and prevent it from restarting. If you stop `vxattachd`, the automatic plex reattachment also stops. To kill the daemon, run the following command from the command line:

```
# ps -afe
```

Locate the process table entry for `vxattachd`, and kill it by specifying its process ID:

```
# kill -9 PID
```

If there is no entry in the process table for `vxattachd`, the automatic site reattachment feature is disabled.

To prevent the automatic site reattachment feature from being restarted, comment out the line that starts `vxattachd` in the `/lib/svc/method/vxvm-recover` startup script and run the `svcadm refresh vxvm/vxvm-recover` command (for Solaris 10 onward), or `/etc/init.d/vxvm-recover` (for OS releases before Solaris 10).

Optimizing I/O performance

- [Chapter 17. Veritas File System I/O](#)
- [Chapter 18. Veritas Volume Manager I/O](#)

Veritas File System I/O

This chapter includes the following topics:

- [About Veritas File System I/O](#)
- [Buffered and Direct I/O](#)
- [Concurrent I/O](#)
- [Cache advisories](#)
- [Freezing and thawing a file system](#)
- [Getting the I/O size](#)
- [About Storage Foundation and High Availability Solutions products database accelerators](#)

About Veritas File System I/O

VxFS processes two basic types of file system I/O:

- Sequential
- Random or I/O that is not sequential

For sequential I/O, VxFS employs a read-ahead policy by default when the application is reading data. For writing, it allocates contiguous blocks if possible. In most cases, VxFS handles I/O that is sequential through buffered I/O. VxFS handles random or nonsequential I/O using direct I/O without buffering.

VxFS provides a set of I/O cache advisories for use when accessing files.

See the *Veritas File System Programmer's Reference Guide*.

See the `vxfsio(7)` manual page.

Buffered and Direct I/O

VxFS responds with read-ahead for sequential read I/O. This results in buffered I/O. The data is prefetched and retained in buffers for the application. The data buffers are commonly referred to as VxFS buffer cache. This is the default VxFS behavior.

On the other hand, direct I/O does not buffer the data when the I/O to the underlying device is completed. This saves system resources like memory and CPU usage. Direct I/O is possible only when alignment and sizing criteria are satisfied.

See [“Direct I/O requirements”](#) on page 486.

All of the supported platforms have a VxFS buffered cache. Each platform also has either a page cache or its own buffer cache. These caches are commonly known as the file system caches.

Direct I/O does not use these caches. The memory used for direct I/O is discarded after the I/O is complete,

Direct I/O

Direct I/O is an unbuffered form of I/O. If the `VX_DIRECT` advisory is set, the user is requesting direct data transfer between the disk and the user-supplied buffer for reads and writes. This bypasses the kernel buffering of data, and reduces the CPU overhead associated with I/O by eliminating the data copy between the kernel buffer and the user's buffer. This also avoids taking up space in the buffer cache that might be better used for something else. The direct I/O feature can provide significant performance gains for some applications.

The direct I/O and `VX_DIRECT` advisories are maintained on a per-file-descriptor basis.

Direct I/O requirements

For an I/O operation to be performed as direct I/O, it must meet certain alignment criteria. The alignment constraints are usually determined by the disk driver, the disk controller, and the system memory management hardware and software.

The requirements for direct I/O are as follows:

- The starting file offset must be aligned to a 512-byte boundary.
- The ending file offset must be aligned to a 512-byte boundary, or the length must be a multiple of 512 bytes.
- The memory buffer must start on an 8-byte boundary.

Direct I/O versus synchronous I/O

Because direct I/O maintains the same data integrity as synchronous I/O, it can be used in many applications that currently use synchronous I/O. If a direct I/O request does not allocate storage or extend the file, the inode is not immediately written.

Direct I/O CPU overhead

The CPU cost of direct I/O is about the same as a raw disk transfer. For sequential I/O to very large files, using direct I/O with large transfer sizes can provide the same speed as buffered I/O with much less CPU overhead.

If the file is being extended or storage is being allocated, direct I/O must write the inode change before returning to the application. This eliminates some of the performance advantages of direct I/O.

Discovered Direct I/O

Discovered Direct I/O is a file system tunable that is set using the `vxtunefs` command. When the file system gets an I/O request larger than the `discovered_direct_iosz`, it tries to use direct I/O on the request. For large I/O sizes, Discovered Direct I/O can perform much better than buffered I/O.

Discovered Direct I/O behavior is similar to direct I/O and has the same alignment constraints, except writes that allocate storage or extend the file size do not require writing the inode changes before returning to the application.

Unbuffered I/O

If the `VX_UNBUFFERED` advisory is set, I/O behavior is the same as direct I/O with the `VX_DIRECT` advisory set, so the alignment constraints that apply to direct I/O also apply to unbuffered I/O. For unbuffered I/O, however, if the file is being extended, or storage is being allocated to the file, inode changes are not updated synchronously before the write returns to the user. The `VX_UNBUFFERED` advisory is maintained on a per-file-descriptor basis.

Data synchronous I/O

If the `VX_DSYNC` advisory is set, the user is requesting data synchronous I/O. In synchronous I/O, the data is written, and the inode is written with updated times and, if necessary, an increased file size. In data synchronous I/O, the data is transferred to disk synchronously before the write returns to the user. If the file is not extended by the write, the times are updated in memory, and the call returns to the user. If the file is extended by the operation, the inode is written before the write returns.

The direct I/O and `VX_DSYNC` advisories are maintained on a per-file-descriptor basis.

Data synchronous I/O vs. synchronous I/O

Like direct I/O, the data synchronous I/O feature can provide significant application performance gains. Because data synchronous I/O maintains the same data integrity as synchronous I/O, it can be used in many applications that currently use synchronous I/O. If the data synchronous I/O does not allocate storage or extend the file, the inode is not immediately written. The data synchronous I/O does not have any alignment constraints, so applications that find it difficult to meet the alignment constraints of direct I/O should use data synchronous I/O.

If the file is being extended or storage is allocated, data synchronous I/O must write the inode change before returning to the application. This case eliminates the performance advantage of data synchronous I/O.

Concurrent I/O

Concurrent I/O (`VX_CONCURRENT`) allows multiple processes to read from or write to the same file without blocking other `read(2)` or `write(2)` calls. POSIX semantics requires `read` and `write` calls to be serialized on a file with other `read` and `write` calls. With POSIX semantics, a `read` call either reads the data before or after the `write` call occurred. With the `VX_CONCURRENT` advisory set, the `read` and `write` operations are not serialized as in the case of a character device. This advisory is generally used by applications that require high performance for accessing data and do not perform overlapping writes to the same file. It is the responsibility of the application or the running threads to coordinate the `write` activities to the same file when using Concurrent I/O.

Concurrent I/O can be enabled in the following ways:

- By specifying the `VX_CONCURRENT` advisory flag for the file descriptor in the `VX_SETCACHE` `ioctl` command. Only the `read(2)` and `write(2)` calls occurring through this file descriptor use concurrent I/O. The `read` and `write` operations occurring through other file descriptors for the same file will still follow the POSIX semantics.
See `vxfsio(7)` manual page.
- By using the `cio` mount option. The `read(2)` and `write(2)` operations occurring on all of the files in this particular file system will use concurrent I/O.
See [“cio mount option”](#) on page 223.
See the `mount_vxfs(1M)` manual page.

Cache advisories

VxFS allows an application to set cache advisories for use when accessing files. VxFS cache advisories enable applications to help monitor the buffer cache and provide information on how better to tune the buffer cache to improve performance gain.

The basic function of the cache advisory is to let you know whether you could have avoided a later re-read of block X if the buffer cache had been a little larger. Conversely, the cache advisory can also let you know that you could safely reduce the buffer cache size without putting block X into jeopardy.

These advisories are in memory only and do not persist across reboots. Some advisories are currently maintained on a per-file, not a per-file-descriptor, basis. Only one set of advisories can be in effect for all accesses to the file. If two conflicting applications set different advisories, both must use the advisories that were last set.

All advisories are set using the `VX_SETCACHE` ioctl command. The current set of advisories can be obtained with the `VX_GETCACHE` ioctl command.

See the `vxfsio(7)` manual page.

Freezing and thawing a file system

Freezing a file system is a necessary step for obtaining a stable and consistent image of the file system at the volume level. Consistent volume-level file system images can be obtained and used with a file system snapshot tool. The freeze operation flushes all buffers and pages in the file system cache that contain dirty metadata and user data. The operation then suspends any new activity on the file system until the file system is thawed.

The `VX_FREEZE` ioctl command is used to freeze a file system. Freezing a file system temporarily blocks all I/O operations to a file system and then performs a sync on the file system. When the `VX_FREEZE` ioctl is issued, all access to the file system is blocked at the system call level. Current operations are completed and the file system is synchronized to disk.

When the file system is frozen, any attempt to use the frozen file system, except for a `VX_THAW` ioctl command, is blocked until a process executes the `VX_THAW` ioctl command or the time-out on the freeze expires.

Getting the I/O size

VxFS provides the `VX_GET_IOPARAMETERS` ioctl to get the recommended I/O sizes to use on a file system. This ioctl can be used by the application to make decisions about the I/O sizes issued to VxFS for a file or file device.

See the `vxtunefs(1M)` and `vxfsio(7)` manual pages.

About Storage Foundation and High Availability Solutions products database accelerators

The major concern in any environment is maintaining respectable performance or meeting performance service level agreements (SLAs). Veritas Storage Foundation products improve the overall performance of database environments in a variety of ways.

These database accelerator technologies enable database performance equal to raw disk partitions, but with the manageability benefits of a file system. With the Dynamic Multi-pathing (DMP) feature of Storage Foundation, performance is maximized by load-balancing I/O activity across all available paths from server to array. DMP supports all major hardware RAID vendors, hence there is no need for third-party multi-pathing software, reducing the total cost of ownership.

Storage Foundation database accelerators enable you to manage performance for your database with more precision.

Veritas Volume Manager I/O

This chapter includes the following topics:

- [Veritas Volume Manager throttling of administrative I/O](#)

Veritas Volume Manager throttling of administrative I/O

In this release, Veritas Volume Manager (VxVM) provides throttling of administrative I/O. During heavy I/O loads, VxVM throttles I/O that it creates to do administrative operations. This behavior ensures that the administrative I/Os do not affect the application I/O performance. When the application I/O load is lighter, VxVM increases the bandwidth usage for administrative I/O operations.

VxVM automatically manages the I/O throttling for administrative tasks, based on its perceived load on the storage. Currently, I/O throttling is supported for the copy operations which use `ATOMIC_COPY` and involve one destination mirror. The I/O throttling is transparent, and does not change the command usage or output. The following commands are supported:

- `vxassist mirror`
- `vxassist snapcreate`
- `vxevac`
- `vxplex att`
- `vxplex cp`
- `vxplex mv`
- `vxprint`
- `vxsnap addmir`

- `vxsnap reattach`
- `vxsd mv`
- `vxtune`

The administrative I/O operations allocate memory for I/O from a separate memory pool. You can tune the maximum size of this pool with the tunable parameter, `vol_max_adminio_poolsz`.

Veritas Extension for Oracle Disk Manager

- [Chapter 19. Using Veritas Extension for Oracle Disk Manager](#)

Using Veritas Extension for Oracle Disk Manager

This chapter includes the following topics:

- [About Oracle Disk Manager](#)
- [About Oracle Disk Manager and Veritas Storage Foundation Cluster File System High Availability](#)
- [About Oracle Disk Manager and Oracle Managed Files](#)
- [Setting up Veritas Extension for Oracle Disk Manager](#)
- [Preparing existing database storage for Oracle Disk Manager](#)
- [Converting Quick I/O files to Oracle Disk Manager files](#)
- [Verifying that Oracle Disk Manager is configured](#)
- [Disabling the Oracle Disk Manager feature](#)
- [Using Cached ODM](#)

About Oracle Disk Manager

Veritas Extension for Oracle Disk Manager is specifically designed for Oracle10g or later to enhance file management and disk I/O throughput. The features of Oracle Disk Manager are best suited for databases that reside in a file system contained in Veritas File System. Oracle Disk Manager allows Oracle10g or later users to improve database throughput for I/O intensive workloads with special I/O optimization.

Veritas Extension for Oracle Disk Manager supports Oracle Resilvering. With Oracle Resilvering, the storage layer receives information from the Oracle database as to which regions or blocks of a mirrored datafile to resync after a system crash. Oracle Resilvering avoids overhead from the VxVM DRL, which increases performance.

Oracle Disk Manager reduces administrative overhead by providing enhanced support for Oracle Managed Files. Veritas Extension for Oracle Disk Manager has Quick I/O-like capabilities, but is transparent to the user. Unlike Veritas Quick I/O, files managed using Veritas Extension for Oracle Disk Manager do not require special file naming conventions. The Oracle Disk Manager interface uses regular database files. If you are upgrading to Oracle10g or later, you should convert from Quick I/O to Oracle Disk Manager.

Database administrators can choose the datafile type used with the Oracle product. Historically, choosing between file system files and raw devices was based on manageability and performance. The exception to this is a database intended for use with Oracle Parallel Server, which requires raw devices on most platforms. If performance is not as important as administrative ease, file system files are typically the preferred file type. However, while an application may not have substantial I/O requirements when it is first implemented, I/O requirements may change. If an application becomes dependent upon I/O throughput, converting data files from file system to raw devices is often necessary.

Oracle Disk Manager was designed to work with Oracle10g or later to provide both performance and manageability. Oracle Disk Manager provides support for Oracle's file management and I/O calls for database storage on VxFS file systems and on raw volumes or partitions. This feature is provided as a dynamically-loaded shared library with which Oracle binds when it is loaded. The Oracle Disk Manager library works with an Oracle Disk Manager driver that is loaded in the kernel to perform its functions.

If you are upgrading to Oracle10g or later, you should convert from Quick I/O to Oracle Disk Manager.

The benefits of using Oracle Disk Manager are as follows:

- True kernel asynchronous I/O for files and raw devices
- Reduced system call overhead
- Improved file system layout by preallocating contiguous files on a VxFS file system
- Performance on file system files that is equal to raw devices
- Transparent to users
- Contiguous datafile allocation

How Oracle Disk Manager improves database performance

Oracle Disk Manager improves database I/O performance to VxFS file systems by:

- Supporting kernel asynchronous I/O
[About kernel asynchronous I/O support](#)
- Supporting direct I/O and avoiding double buffering
[About direct I/O support and avoiding double buffering](#)
- Avoiding kernel write locks on database files
[About avoiding kernel write locks on database files](#)
- Supporting many concurrent I/Os in one system call
[About supporting many concurrent I/Os in one system call](#)
- Avoiding duplicate opening of files per Oracle instance
[About avoiding duplicate file open calls](#)
- Allocating contiguous data files
[About allocating contiguous data files](#)

About kernel asynchronous I/O support

Asynchronous I/O performs non-blocking system level reads and writes, allowing the system to perform multiple I/O requests simultaneously. Kernel asynchronous I/O is better than library asynchronous I/O because the I/O is queued to the disk device drivers in the kernel, minimizing context switches to accomplish the work.

About direct I/O support and avoiding double buffering

I/O on files using `read()` and `write()` system calls typically results in data being copied twice: once between the user and kernel space, and the other between kernel space and the disk. In contrast, I/O on raw devices is copied directly between user space and disk, saving one level of copying. As with I/O on raw devices, Oracle Disk Manager I/O avoids the extra copying. Oracle Disk Manager bypasses the system cache and accesses the files with the same efficiency as raw devices. Avoiding double buffering reduces the memory overhead on the system. Eliminating the copies from kernel to user address space significantly reduces kernel mode processor utilization, freeing more processor cycles to execute the application code.

About avoiding kernel write locks on database files

When database I/O is performed by way of the `write()` system call, each system call acquires and releases a kernel write lock on the file. This lock prevents simultaneous write operations on the same file. Because database systems usually implement their own locks for managing concurrent access to files, write locks

unnecessarily serialize I/O writes. Oracle Disk Manager bypasses file system locking and lets the database server control data access.

About supporting many concurrent I/Os in one system call

When performing asynchronous I/O, an Oracle process may try to issue additional I/O requests while collecting completed I/Os, or it may try to wait for particular I/O requests synchronously, as it can do no other work until the I/O is completed. The Oracle process may also try to issue requests to different files. All this activity can be accomplished with one system call when Oracle uses the Oracle Disk Manager I/O interface. This interface reduces the number of system calls performed to accomplish the same work, reducing the number of user space/kernel space context switches.

About avoiding duplicate file open calls

Oracle Disk Manager allows files to be opened once, providing a “file identifier.” This is called “identifying” the files. The same file identifiers can be used by any other processes in the Oracle instance. The file status is maintained by the Oracle Disk Manager driver in the kernel. The reduction in file open calls reduces processing overhead at process initialization and termination, and it reduces the number of file status structures required in the kernel.

About allocating contiguous data files

Oracle Disk Manager can improve performance for queries, such as sort and parallel queries, that use temporary tablespaces. Without Oracle Disk Manager, Oracle does not initialize the data files for the temporary tablespaces. Therefore, the data files become sparse files and are generally fragmented. Sparse or fragmented files lead to poor query performance. When using Oracle Disk Manager, the data files are initialized for the temporary tablespaces and are allocated in a contiguous fashion, so that they are not sparse.

About Oracle Disk Manager and Veritas Storage Foundation Cluster File System High Availability

Oracle Disk Manager (ODM) supports access to clustered files in the SFCFSHA environment. With a Veritas Storage Foundation Cluster File System High Availability license, ODM supports SFCFSHA files in a serially-exclusive mode which allows access to each SFCFSHA file by one node at a time, but does not allow simultaneous access from multiple nodes.

See the `mount_odm(1)` man page for more information on its cluster support modes.

About Oracle Disk Manager and Oracle Managed Files

Oracle10g or later offers a feature known as Oracle Managed Files (OMF). OMF manages datafile attributes such as file names, file location, storage attributes, and whether or not the file is in use by the database. OMF is only supported for databases that reside in file systems. OMF functionality is greatly enhanced by Oracle Disk Manager.

OMF is a file management feature that:

- Eliminates the task of providing unique file names
- Offers dynamic space management by way of the tablespace auto-extend functionality of Oracle10g or later

The main requirement for OMF is that the database be placed in file system files. There are additional prerequisites imposed upon the file system itself.

OMF should only be used in file systems that reside within striped logical volumes, which support dynamic file system growth. File systems intended for OMF use must also support large, extensible files in order to facilitate tablespace auto-extension. Raw partitions cannot be used for OMF.

By default, OMF data files are created with auto-extensibility. This attribute reduces capacity planning associated with maintaining existing databases and implementing new applications. Due to disk fragmentation that occurs as the tablespace grows over time, database administrators have been somewhat cautious when considering auto-extensible tablespaces. Oracle Disk Manager eliminates this concern.

When Oracle Disk Manager is used in conjunction with OMF, special care is given within Veritas Extension for Disk Manager to ensure that contiguous disk space is allocated to data files, including space allocated to a tablespace when it is auto-extended. The table and index scan throughput does not decay as the tablespace grows.

How Oracle Disk Manager works with Oracle Managed Files

The following example illustrates the relationship between Oracle Disk Manager and Oracle Managed Files (OMF). The example shows the `init.ora` contents and the command for starting the database instance. To simplify Oracle UNDO management, the new Oracle10g or later `init.ora` parameter `UNDO_MANAGEMENT` is set to `AUTO`. This is known as System-Managed Undo.

Note: Before building an OMF database, you need the appropriate `init.ora` default values. These values control the location of the `SYSTEM` tablespace, online redo logs, and control files after the `CREATE DATABASE` statement is executed.

```
$ cat initPROD.ora
UNDO_MANAGEMENT = AUTO
DB_CREATE_FILE_DEST = '/PROD'
DB_CREATE_ONLINE_LOG_DEST_1 = '/PROD'
db_block_size = 4096
db_name = PROD
$ sqlplus /nolog
SQL> connect / as sysdba
SQL> startup nomount pfile= initPROD.ora
```

The Oracle instance starts.

```
Total System Global Area 93094616 bytes
Fixed Size 279256 bytes
Variable Size 41943040 bytes
Database Buffers 50331648 bytes
Redo Buffers 540672 bytes
```

To implement a layout that places files associated with the `EMP_TABLE` tablespace in a directory separate from the `EMP_INDEX` tablespace, use the `ALTER SYSTEM` statement. This example shows how OMF handles file names and storage clauses and paths. The layout allows you to think of the tablespaces as objects in a file system as opposed to a collection of data files. Since OMF uses the Oracle Disk Manager file resize function, the tablespace files are initially created with the default size of 100 MB and ODM increases the size as needed. Use the `MAXSIZE` attribute to limit growth.

The following example shows the commands for creating an OMF database and for creating the `EMP_TABLE` and `EMP_INDEX` tablespaces in their own locale:

```
SQL> create database PROD;
```

Note: The directory must exist for OMF to work, so the `SQL*Plus HOST` command is used to create the directories:

The database is created.

```
SQL> HOST mkdir /PROD/EMP_TABLE;
SQL> ALTER SYSTEM SET DB_CREATE_FILE_DEST = '/PROD/EMP_TABLE';
```

The system is altered.

```
SQL> create tablespace EMP_TABLE DATAFILE AUTOEXTEND ON MAXSIZE \
500M;
```

A tablespace is created.

```
SQL> ALTER SYSTEM SET DB_CREATE_FILE_DEST = '/PROD/EMP_INDEX';
```

The system is altered.

```
SQL> create tablespace EMP_INDEX DATAFILE AUTOEXTEND ON MAXSIZE \
100M;
```

A tablespace is created.

Use the `ls` command to show the newly created database:

```
$ ls -lFR
total 638062
drwxr-xr-x 2 oracle10g dba 96 May  3 15:43 EMP_INDEX/
drwxr-xr-x 2 oracle10g dba 96 May  3 15:43 EMP_TABLE/
-rw-r--r-- 1 oracle10g dba 104858112 May  3 17:28 ora_1_BEhYgc0m.log
-rw-r--r-- 1 oracle10g dba 104858112 May  3 17:27 ora_2_BEhYu4NA.log
-rw-r--r-- 1 oracle10g dba 806912 May  3 15:43 ora_BEahlfUX.ctl
-rw-r--r-- 1 oracle10g dba 10489856 May  3 15:43 ora_sys_undo_BEajPSVq.dbf
-rw-r--r-- 1 oracle10g dba 104861696 May  3 15:4 ora_system_BEaiFE8v.dbf
-rw-r--r-- 1 oracle10g dba 186 May  3 15:03 PROD.ora

./EMP_INDEX:
total 204808
-rw-r--r-- 1 oracle10g dba 104861696 May  3 15:43
ora_emp_inde_BEakGfun.dbf

./EMP_TABLE:
total 204808
-rw-r--r-- 1 oracle10g dba 104861696 May  3 15:43
ora_emp_tabl_BEak1LqK.dbf
```

Setting up Veritas Extension for Oracle Disk Manager

Veritas Extension for Oracle Disk Manager is part of Veritas Storage Foundation Cluster File System High Availability. Veritas Extension for Oracle Disk Manager is enabled once Veritas Storage Foundation Cluster File System High Availability and Oracle10g or later are installed. The Veritas Extension for Oracle Disk Manager library is linked to the library in the `{ORACLE_HOME}/lib` directory.

If you are performing a local Oracle installation, not on the Cluster file system, then ODM linking needs to be performed on all nodes in the cluster.

Before setting up Veritas Extension for Oracle Disk Manager, the following conditions must be met:

Prerequisites	■ Veritas Storage Foundation Cluster File System High Availability must be installed on your system. ■ Oracle10g, or later, must be installed on your system. ■ If Cached Quick I/O is available, do not enable Oracle Disk Manager when Cached Quick I/O is enabled for data files.
Usage Notes	■ When the Quick I/O feature is available, Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O. Do not turn off the Quick I/O mount option, which is the default. ■ Oracle uses default file access methods if Oracle10g or later, Veritas Storage Foundation Cluster File System High Availability is not installed, or VxFS 5.1 SP1 is not available in the kernel.

Preparing existing database storage for Oracle Disk Manager

Non-Quick I/O files in a VxFS file system work with Oracle Disk Manager without any changes. The files are found and identified for Oracle Disk Manager I/O by default. To take full advantage of Oracle Disk Manager data files, files should not be fragmented.

If you are using Quick I/O files in a VxFS file system and you want to move to Oracle Disk Manager, convert the Quick I/O files to normal files using the `qio_convertdbfiles -u` command.

You must be running Oracle10g or later to use Oracle Disk Manager.

Converting Quick I/O files to Oracle Disk Manager files

If you plan to run the Veritas product with Oracle10g or later, and you have been using Quick I/O files, Symantec recommends that you convert your Quick I/O files to regular files. This should be done after you upgrade.

Note: If you are running an earlier version of Oracle (Oracle 8.x or lower), you should not convert your Quick I/O files because Oracle Disk Manager is for Oracle10g or later only.

The Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option, which is the default.

To convert Quick I/O files to Oracle Disk Manager files

- 1 As Oracle DBA, run `qio_getdbfiles` to retrieve a list of all data files.

```
$ /opt/VRTS/bin/qio_getdbfiles -T ora -a
```

The list is compiled in a file named `mkqio.dat`.

- 2 Shutdown the database.
- 3 As Oracle DBA, run `qio_convertdbfiles` in the directory containing the `mkqio.dat` file. The `qio_convertdbfiles` script converts all Quick I/O files to ODM files.

```
$ /opt/VRTS/bin/qio_convertdbfiles -T ora -u
```

- 4 Restart the database instance.

Verifying that Oracle Disk Manager is configured

Before verifying that Oracle Disk Manager is configured, make sure that the following conditions are met:

Prerequisites

- `/opt/VRTSodm/lib/libodm.so` must exist.
- If you are using Oracle 10g, `$ORACLE_HOME/lib/libodm10.so` is linked to `/opt/VRTSodm/lib/sparcv9/libodm.so`.
- If you are using Oracle 11g, `$ORACLE_HOME/lib/libodm11.so` is linked to `/opt/VRTSodm/lib/sparcv9/libodm.so`.
- If you are using Oracle10g on x64 Operating System, `$ORACLE_HOME/lib/libodm10.so` is linked to `/opt/VRTSodm/lib/amd64/libodm.so`.

To verify that Oracle Disk Manager is configured

- 1 Verify that the ODM feature is included in the license:

```
# /opt/VRTS/bin/vxlicrep | grep ODM
```

The output verifies that ODM is enabled.

Note: Verify that the license key containing the ODM feature is not expired. If the license key has expired, you will not be able to use the ODM feature.

- 2 Check that the `VRTSodm` package is installed:

```
# pkginfo -l VRTSodm
system VRTSodm Veritas Oracle Disk Manager
```

- 3 Check that `libodm.so` is present.

If you are running 32-bit Oracle9i, use the following command:

```
# ls -lL /opt/VRTSodm/lib/libodm.so
-rw-r--r-- 1 root sys 14336 Apr 25 18:42
/opt/VRTSodm/lib/libodm.so
```

If you are running 64-bit Oracle9i, use the following command:

```
# ls -lL /opt/VRTSodm/lib/sparcv9/libodm.so
-rw-r--r-- 1 root sys 14336 Apr 25 18:42
/opt/VRTSodm/lib/sparcv9/libodm.so
```

To verify that Oracle Disk Manager is running

- 1 Start the Oracle database.
- 2 Check that the instance is using the Oracle Disk Manager function:

```
# cat /dev/odm/stats
# echo $?
0
```

- 3 Verify that the Oracle Disk Manager is loaded:

```
# modinfo | grep ODM | grep VRTS
162 7b76c000 184a0 25 1 odm (VRTS ODM 5.1.10.00,REV=MPlu)
```

- 4 In the alert log, verify the Oracle instance is running. The log should contain output similar to the following:

```
Oracle instance running with ODM: Veritas 5.1.100.00 ODM Library,
Version 2.0
```

Disabling the Oracle Disk Manager feature

Since the Oracle Disk Manager feature uses regular files, you can access these files as regular VxFS files as soon as the feature is disabled.

The steps for disabling the Oracle Disk Manager feature are the same for both 32- and 64-bit Oracle10g.

Note: To convert to VxFS with Quick I/O, disable Oracle Disk Manager using the following procedure, then convert the files to Quick I/O files.

See [“Converting Quick I/O files to Oracle Disk Manager files”](#) on page 501.

Before disabling the Oracle Disk Manager feature, you may want to back up your files.

To disable the Oracle Disk Manager feature in an Oracle instance

- 1 Shut down the database instance.
- 2 Use the `rm` and `ln` commands to remove the link to the Oracle Disk Manager Library.

For Oracle 11g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm11.so
# ln -s ${ORACLE_HOME}/lib/libodmd11.so \
${ORACLE_HOME}/lib/libodm11.so
```

For Oracle 10g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm10.so
# ln -s ${ORACLE_HOME}/lib/libodmd10.so \
${ORACLE_HOME}/lib/libodm10.so
```

For Oracle 10g on Opteron, enter:

```
# rm ${ORACLE_HOME}/lib/libodm10.so
# ln -s ${ORACLE_HOME}/lib/libodmd10.so \
${ORACLE_HOME}/lib/libodm10.so
```

- 3 Restart the database instance.

Using Cached ODM

ODM I/O normally bypasses the file system cache and directly reads from and writes to disk. Cached ODM enables some I/O to use caching and read ahead, which can improve ODM I/O performance. Cached ODM performs a conditional form of caching that is based on per-I/O hints from Oracle. The hints indicate what Oracle does with the data. ODM uses these hints to perform caching and read ahead for some reads, but ODM avoids caching other reads, even for the same file.

You can enable cached ODM for local mount files and cluster mount files.

See [“Enabling Cached ODM for file systems”](#) on page 506.

Cached ODM can be configured in two ways. The primary configuration method is to turn caching on or off for all I/O on a per-file basis. The secondary configuration method is to adjust the ODM cachemap. The cachemap maps file type and I/O type combinations into caching advisories.

See [“Modifying Cached ODM settings for individual files”](#) on page 506.

See [“Adding Cached ODM settings via the cachemap”](#) on page 507.

Enabling Cached ODM for file systems

Cached ODM is initially disabled on a file system. You enable Cached ODM for a file system by setting the `odm_cache_enable` option of the `vxtunefs` command after the file system is mounted.

See the `vxtunefs(1M)` manual page.

Note: The `vxtunefs` command enables conditional caching for all of the ODM files on the file system.

To enable Cached ODM for a file system

- 1 Enable Cached ODM on the VxFS file system `/database01`:

```
# vxtunefs -s -o odm_cache_enable=1 /database01
```

- 2 Optionally, you can make this setting persistent across mounts by adding a file system entry in the file `/etc/vx/tunefstab`:

```
/dev/vx/dsk/datadg/database01 odm_cache_enable=1
```

See the `tunefstab(4)` manual page.

Modifying Cached ODM settings for individual files

You can use the `odmadm setcachefile` command to override the cachemap for a specific file so that ODM caches either all or none of the I/O to the file. The caching state can be ON, OFF, or DEF (default). The DEF caching state is conditional caching, meaning that for each I/O, ODM consults the cachemap and determines whether the specified file type and I/O type combination should be cached. The ON caching state causes the specified file always to be cached, while the OFF caching state causes the specified file never to be cached.

See the `odmadm(1M)` manual page.

Note: The cache advisories operate only if Cached ODM is enabled for the file system. If the `odm_cache_enable` flag is zero, Cached ODM is OFF for all of the files in that file system, even if the individual file cache advisory for a file is ON.

To enable unconditional caching on a file

- ◆ Enable unconditional caching on the file `/mnt1/file1`:

```
# odmadm setcachefile /mnt1/file1=on
```

With this command, ODM caches all reads from `file1`.

To disable caching on a file

- ◆ Disable caching on the file `/mnt1/file1`:

```
# odmadm setcachefile /mnt1/file1=off
```

With this command, ODM does not cache reads from `file1`.

To check on the current cache advisory settings for a file

- ◆ Check the current cache advisory settings of the files `/mnt1/file1` and `/mnt2/file2`:

```
# odmadm getcachefile /mnt1/file1 /mnt2/file2
/mnt1/file1,ON
/mnt2/file2,OFF
```

To reset all files to the default cache advisory

- ◆ Reset all files to the default cache advisory:

```
# odmadm resetcachefiles
```

Adding Cached ODM settings via the cachemap

You can use the `odmadm setcachemap` command to configure the cachemap. The cachemap maps file type and I/O type combinations to caching advisories. ODM uses the cachemap for all files that have the default conditional cache setting. Such files are those for which caching has not been turned on or off by the `odmadm setcachefile` command.

See the `odmadm(1M)` manual page.

By default, the cachemap is empty, but you can add caching advisories by using the `odmadm setcachemap` command.

To add caching advisories to the cachemap

- ◆ Add a caching advisory to the cachemap:

```
# odmadm setcachemap data/data_read_seq=cache,readahead
```

With this example command, ODM uses caching and readahead for I/O to online log files (`data`) that have the `data_read_seq` I/O type. You can view the valid file type and I/O type values from the output of the `odmadm getcachemap` command.

See the `odmadm(1M)` manual page.

Making the caching settings persistent across mounts

By default, the Cached ODM settings are not persistent across mounts. You can make the settings persistent by creating the `/etc/vx/odmadm` file and listing the caching advisory settings in the file

To make the caching setting persistent across mounts

- ◆ Create the `/etc/vx/odmadm` file to list files and their caching advisories. In the following example of the `/etc/vx/odmadm` file, if you mount the `/dev/vx/dsk/rootdg/voll` device at `/mnt1`, `odmadm` turns off caching for `/mnt1/oradata/file1`:

```
setcachemap data/read_data_header=cache
setcachemap all/datapump=cache,readahead
device /dev/vx/dsk/rootdg/voll
setcachefile oradata/file1=off
```


Clustered NFS

- [Chapter 20. Clustered NFS](#)

Clustered NFS

This chapter includes the following topics:

- [About Clustered NFS](#)
- [Requirements for Clustered NFS](#)
- [Understanding how Clustered NFS works](#)
- [cfsshare manual page](#)
- [Configure and unconfigure Clustered NFS](#)
- [Reconciling major and minor numbers for NFS shared disks](#)
- [Administering Clustered NFS](#)
- [How to mount an NFS-exported file system on the NFS clients](#)
- [Disabling SMF for NFS daemons on Solaris 10](#)
- [Debugging Clustered NFS](#)

About Clustered NFS

The Clustered NFS (CNFS) feature gracefully handles failure of any node and reclaim the locks in such a way as to not accidentally lose any existing lock grants without notification.

This release only supports NFS Version 3.

See [“How to mount an NFS-exported file system on the NFS clients”](#) on page 532.

Requirements for Clustered NFS

- Prior knowledge of NFS and NFS locking is a prerequisite.

Understanding how Clustered NFS works

This Clustered NFS feature allows the same file system mounted across multiple nodes using CFS to be shared over NFS from any combination of those nodes without any loss of functionality during failover. The failover of NFS lock servers includes all the locks being released by the old node then reclaimed by clients talking to the new node during the grace period.

Basic design

The basic design is to have VCS manage virtual IP (VIP) resources that can failover between nodes and to add extra code into the steps used to handle these resources to properly handle the NFS level operations. All other involved resources are active on all nodes participating. The lock data, which is saved into a shared area, is managed and used in lock-step with the virtual IP resources to ensure that all locks are reclaimed properly by clients while preventing any inappropriate locks from being taken at the wrong time. Such interfering locks are prevented through a combination of stopping services and using new features of the lock handling inside the VxFS kernel code.

To communicate with the new code in the kernel, the `fsclustadm` command has been modified to add command line interfaces to the private ioctl calls.

Note: You must have at least one VIP configured per each CNFS server.

Internal Clustered NFS functionality

This section describes the internal functionality of the triggers and actions scripts that are a part of the Clustered NFS solution.

preonline trigger

The `preonline` script copies the lock state files created by the node status monitor (normally called `statd` or `rpc.statd` daemon) during IP failovers and node failures.

The `preonline` script does the following on IP failover or node failure:

- Finds the IP and the node it was last online on.
- Finds the node on which the IP is next going to go online on.
- Check if the `/locks/sm/lastonline/sm/` directory is empty. If it is empty, then exit as there is no lock. There is no need to restart lockd.

- Calls `/opt/VRTS/bin/fsclustadm frlpause_enable` and `/opt/VRTS/bin/fsclustadm frlock_pause` to ensure that file system does not give out any new locks during the failover.
- Stops lock and status services on all nodes to prevent granting locks.
- Copies all the files from `/locks/sm/lastonline/sm/` to `/locks/sm/nextonline/sm/` directory.
 where *locks* is the file system created for storing lock information.
 where *lastonline* is the node on which the VIP resource was previous online.
 where *nextonline* is the node on which the VIP resource will go online next.
- Calls `/opt/VRTS/bin/fsclustadm frlock_resume` to resume giving out locks.

Note: At the end of the preonline trigger all lock services have stopped on all nodes and no NFS locks can be requested until they are restarted.

postonline trigger

The `postonline` script for each VIP does the following during an IP failover or node failure:

- Check if the `/locks/sm/lastonline/sm/` directory is empty. If it is empty, then there is no need to restart the lock service.
- Starts lock services, triggers reclaim, and grace mode on all nodes.
- The restarting of status monitor scans all lock status files in the state directory and contacts all nodes to reclaim their locks. The state files get deleted after they are processed and reclaim messages sent appropriately.
- The lock server goes into grace mode and only allows clients to recover their locks during the grace period. It does not give out any new locks during the grace period.

postoffline trigger

The `postoffline` script does the following on IP failover:

- Calls `/opt/VRTS/bin/fsclustadm frlpause_disable` to reduce the internal usage counter.
- Each call to `/opt/VRTS/bin/fsclustadm frlpause_enable` needs to be matched with a call to `/opt/VRTS/bin/fsclustadm frlpause_disable` as the kernel keeps an internal counter to track the number of IP addresses active on a system. If there are no active IPs on a system, it will be in disabled mode.

Note: This trigger is only called for an administratively initiated failover. An actual system failure and reboot discards the local state being manipulated in this stage. This is the one trigger called on the node that was previously hosting the VIP, while the others are called on the server taking over.

Actions

- On each node, a `/opt/VRTSvcs/bin/IP/actions/nfscfs` file is installed. This file is used to start and stop the NFS locking daemons on a specified node. The `action` script is used instead of using `rsh`, `ssh` or `hacli` for remote command execution from the triggers.
- On each node, a `/opt/VRTSvcs/bin/ApplicationNone/actions/nfscfsapp` file is installed. This file is used while configuring and unconfiguring the Clustered NFS solution using `cfsshare config` and `cfsshare unconfig` commands.

cfsshare manual page

This Clustered NFS feature adds a new configuration utility called `cfsshare` to the `VRTScavf` package and several scripts that are added into the VCS configuration to manage parallel NFS server resources. The `cfsshare` command modifies VCS resources that were created by other utilities such as `cfsmntadm`.

See the `cfsshare(1M)` manual page.

Configure and unconfigure Clustered NFS

This section describes how to configure and unconfigure Clustered NFS.

Configure Clustered NFS

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

The CNFS solution requires a shared file system such as `/locks` that is mounted on all cluster nodes. This file system is not a data file system; it contains the lock state files corresponding to NFS clients holding locks on the CNFS servers.

```
cfsshare config -p nfs [-n] shared_disk_group shared_volume mount_point
```

Note: If CIFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of Clustered NFS.

The local state tracking directory contains a file for each NFS client that has a transaction with the NFS server. The local state tracking directory is:

```
/var/statmon/sm
```

This creates a symlink to `/locks/sm/nodename/sm` on all the cluster nodes. This allows the lock state files for any cluster node to be accessed by other nodes in the cluster, even when the node is down.

The `-n` option can be used if the user does not want `cfsshare` to create the symlink to `/locks/sm/nodename/sm`. If this option is used, then the user needs to manually create a symlink.

The `config` option adds this shared file system to VCS configuration; it creates the corresponding CFSSMount resource in a special parallel service group called `cfsnfssg`. This also creates an NFS resource in the `cfsnfssg` service group. In addition to this, a separate resource of the new type `ApplicationNone` is created to monitor `lockd` and `statd` daemons.

If you run the `cfsshare config -n` option, you need to perform the following procedure:

- 1 On each node, create the following directory inside the locks directory, if it does not already exist:

```
# mkdir -p /locks/sm/nodename/sm
```

- 2 On each cluster node, create a symlink from `/locks/sm/nodename/sm` to `/var/statmon/sm`.

```
# ln -sf /locks/sm/nodename/sm /var/statmon
```

- 3 Run the following commands on any one cluster node to set the owner, group, and permissions of `/locks/sm` appropriately:

```
# chown -R daemon:daemon /locks/sm
# chmod -R 755 /locks/sm
```

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m user -l /var/run \
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```

See the Common Internet File System chapter for explanation regarding the various CIFS-related options.

Service group `cfsnfssg_dummy`

As part of CNFS configuration, a service group called `cfsnfssg_dummy` gets created. This service group is mostly offline.

There is a limit on the number of service groups that can be created in VCS. If this limit is reached, then `cfsnfssg_dummy` serves as the service group in which resources get created during `cfsshare unshare` and `cfsshare delete` operations.

See the *Veritas Cluster Server Administrator's Guide* for information about the `GroupLimit` attribute.

Unconfiguring Clustered NFS

This command is used to undo all the steps during the config phase.

```
cfsshare unconfig -p nfs
```

Note: If there are any CFS file systems still being shared or any virtual IP is added the `cfsshare unconfig` command fails.

Reconciling major and minor numbers for NFS shared disks

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as a Solaris partition or a VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system.

Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Checking the major and minor number for VxVM volumes

The following sections describe checking and changing, if necessary, the major and minor numbers for the VxVM volumes that cluster systems use.

To check major and minor numbers on VxVM volumes

- 1 Place the VCS command directory in your path. For example:

```
# export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
```

- 2 To list the devices, use the `ls -lL block_device` command on each node:

```
# ls -lL /dev/vx/dsk/shareddg/vol3
```

On Node A, the output may resemble:

```
brw----- 1 root root 32,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

On Node B, the output may resemble:

```
brw----- 1 root root 36,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

- 3 Import the associated shared disk group on each node.
- 4 Use the following command on each node exporting an NFS file system. The command displays the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses:

Note: Other major numbers are also displayed, but only `vxio` and `vxspec` are of concern for reconciliation.

```
# grep vx /etc/name_to_major
```

Output on Node A:

```
vxdump 30
vxio 32
vxspec 33
vxfen 87
vxglm 91
```

Output on Node B:

```
vxdump 30
vxio 36
vxspec 37
vxfen 87
vxglm 91
```


- 5 To change Node B's major numbers for `vxio` and `vxspec` to match those of Node A, use the command:

```
haremajor -vx major_number_vxio major_number_vxspec
```

For example:

```
# haremajor -vx 32 33
```

If the command succeeds, proceed to step 8. If this command fails, you receive a report similar to the following:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

- 6 If you receive this report, use the `haremajor` command on Node A to change the major number (32/33) to match that of Node B (36/37). For example, enter:

```
# haremajor -vx 36 37
```

If the command fails again, you receive a report similar to the following:

```
Error: Preexisting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

- 7 If you receive the second report, choose the larger of the two available numbers (in this example, 128). Use this number in the `haremajor` command to reconcile the major numbers. Type the following command on both nodes:

```
# haremajor -vx 128 129
```

- 8 Reboot each node on which `haremajor` was successful.
- 9 If the minor numbers match, proceed to reconcile the major and minor numbers of your next NFS block device.
- 10 If the block device on which the minor number does not match is a volume, consult the `vxldg(1M)` manual page. The manual page provides instructions on reconciling the Veritas Volume Manager minor numbers, and gives specific reference to the `reminor` option.

Nodes where the `vxio` driver number has changed require rebooting.

Administering Clustered NFS

This section describes Clustered NFS scenarios.

See [“Samples for configuring a Clustered NFS”](#) on page 522.

See [“Sample main.cf file”](#) on page 526.

Displaying the NFS shared CFS file systems

This command displays the CFS file systems that are currently being NFS shared by the cluster nodes.

```
cfsshare display
```

Sharing a CFS file system previously added to VCS

Note: You can use the `-N` option to specify the NFS share options.

You can also use the `-p nfs` to specify that the protocol to be used is NFS.

See the `cfsshare(1M)` manual page for more information.

Before running this command, the user should have run `cfsmntadm` command to add the shared file system to VCS configuration and the `cfsmount` command to mount the shared file system at the *mount_point*. Once these commands have been executed, the CFSMount resource corresponding to the *mount_point* gets created in either a default service group (with a name similar to `vrts_vea_cfs_int_cfsmountnumber`) or in a separate service group, as specified by the user.

The `cfsshare share` command moves the CFSMount resource corresponding to the *mount_point* and the associated CVMVolDg resource to the `cfsnfssg` service group (that was created using the `config` option). In addition, this command also creates a share resource on top of the CFSMount resource in the same `cfsnfssg` service group.

```
cfsshare share mount_point [share_options]
```

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSMount and CVMVolDg resources in the `cfsnfssg` service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Unsharing the previous shared CFS file system

Before running this command, the user is supposed to have run the `cfsshare share` command.

The `cfsshare unshare` command enables the user to stop sharing the file system mounted at the *mount_point*. This command moves the Share, CFSSMount, and CVMVolDg resources corresponding to the *mount_point* from cfsnfssg service group to a newly created service group. The Share resource is taken offline and then deleted.

```
cfsshare unshare mount_point
```

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSSMount and CVMVolDg resources in the newly created service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Running the `cfsmntadm delete` command does not remove the `ActivationMode` attribute. If no volumes or vsets in the disk group are in the VCS config, you must use the `cfsdgadm delete` to remove this `ActivationMode` attribute.

Adding an NFS shared CFS file system to VCS

This command adds the CFS file system to the VCS configuration in the cfsnfssg service group, then mounts the file system at the *mount_point*, and NFS shares the CFS file system.

```
cfsshare add [-D] shared_disk_group shared_volume mount_point \  
[share_options] node_name=[mount_options]...
```

```
cfsshare add [-D] shared_disk_group shared_volume mount_point \  
[share_options] all=[mount_options]
```

```
cfsshare add -p nfs [-D] [-N nfs_share_options] shared_disk_group \  
shared_volume mount_point <node_name=[mount_point]
```

```
cfsshare add -p nfs [-D] [-N nfs_share_options] shared_disk_group \  
shared_volume mount_point all=[mount_point]
```

Deleting the NFS shared CFS file system from VCS

Before running this command, the user is supposed to have run the `cfsshare add` command to create the required resources (Share, CFSMount, and CVMVolDg, if needed) in the `cfsnfssg` service group.

This command unshares the CFS file system mounted at the *mount_point*, unmounts the CFS file system, and removes the CFS file system from the VCS configuration.

```
cfsshare delete mount_point
```

Adding a virtual IP address to VCS

This command is used to create a new non-parallel/failover service group that contains a NIC resource for the given network device and an IP resource for the virtual IP address.

```
cfsshare addvip [ -a nodename ] network_interface address netmask
```

The `cfsshare addvip` command lets you specify only one network interface, that is assumed to be present on all cluster nodes. If you want to specify different network interfaces for different cluster nodes, then you need to run certain VCS commands. For example:

```
# haconf -makerw
# hares -local vip1 Device
# hares -modify vip1 Device bge1 -sys system01
# hares -modify vip1 Device bge2 -sys system02
# hares -local nic1 Device
# hares -modify nic1 Device bge1 -sys system01
# hares -modify nic1 Device bge2 -sys system02
# haconf -dump -makero
```

where *vip1* is the virtual IP resource created by the `cfsshare addvip` command.

where *nic1* is the NIC resource respectively created by the `cfsshare addvip` command.

where *system01* and *system02* are the cluster nodes.

Deleting a virtual IP address from VCS

This command is used to delete the non-parallel/failover service group corresponding to the virtual IP address.

```
cfsshare deletevip address
```

Adding an IPv6 virtual IP address to VCS

This command is used to create a new non-parallel/failover service group which contains a NIC resource for the given network device and an IP resource for the IPv6 virtual IP address.

```
cfsshare addvip6 [ -a nodename ] network_interface ipv6_address \  
prefixlen
```

Deleting an IPv6 virtual IP address from VCS

This command is used to delete the non-parallel/failover service group corresponding to the IPv6 virtual IP address.

```
cfsshare deletevip6 ipv6_address
```

Changing the share options associated with an NFS share

This section describes how to change the share options associated with an NFS share.

To change the share options associated with an NFS share

- 1 On any node in the cluster, run `cfsshare unshare` to unshare the file system:

```
# cfsshare unshare mount_point
```

- 2 On any node in the cluster, run `cfsshare share` to share the file system with the desired share options:

```
# cfsshare share -p nfs mount_point share_options
```

Note: The `cfsshare unshare` operation can affect NFS clients that might have mounted the `mount_point` file system.

Sharing a file system checkpoint

This section describes how to share a file system checkpoint.

To share a file system checkpoint

- 1 To add the checkpoint to the VCS configuration, enter:

```
# cfsmntadm add ckpt ckptname mntpt_of_fs mntpt_of_checkpoint \
all=cluster,rw
```

where *ckptname* is the checkpoint name.

where *mntpt_of_fs* is the name of the mount point of the file system.

where *mntpt_of_checkpoint* is the mount point for the checkpoint.

- 2 To mount the checkpoint, enter:

```
# cfsmount mntpt_of_checkpoint
```

- 3 Run the `cfsshare share` command to share this checkpoint:

```
# cfsshare share -p nfs mntpt_of_checkpoint
```

Samples for configuring a Clustered NFS

There are two samples for configuring a Clustered NFS.

Note: Ensure that you have setup a shared disk group with shared volumes and VxFS file system on top.

Sample 1 uses the `cfsshare` command to configure

This sample is intended to use `cfsshare` command to configure and control this feature.

To configure a Clustered NFS (Sample 1)

- 1 Configure a VCS configuration for CFS/CVM, enter:

```
# cfscluster config
```

- 2 Configure CNFS components, enter:

```
# cfsshare config -p nfs shared_disk_group shared_volume \
mount_point
```

For example:

```
# cfsshare config -p nfs cfsdg vollocks /locks
```

- 3 Add and mount the CFS file system to the VCS configuration, enter:

```
# cfsmntadm add [-D] shared_disk_group shared_volume mount_point \  
[service_group] all=[mount_options]  
# cfsmount mount_point
```

For example:

```
# cfsmntadm add cfsdg vol1 /mnt1 all=delaylog,largefiles  
# cfsmount /mnt1
```

- 4 Share the CFS file system, enter:

```
# cfsshare share mount_point [share_options]
```

For example:

```
# cfsshare share /mnt1 rw
```

- 5 Run the following command to display the currently exported file systems:

```
# cfsshare display  
CNFS metadata filesystem : /locks  
Protocols Configured : NFS  
#RESOURCE      MOUNTPOINT      PROTOCOL  OPTIONS  
share1         /defragvol      NFS       rw,no_root_squash
```

- 6 Add the virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvip [ -a nodename ] \  
network_interface address netmask
```

For example:

```
# cfsshare addvip bge0 \  
10.182.111.161 255.255.240.0
```

- 7 Delete a previously added virtual IP address from the configuration, enter:

```
# cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.111.161
```

- 8 Add the IPv6 virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvipv6 [ -a nodename ] network_interface \  
ipv6_address prefixlen
```

For example:

```
# cfsshare addvipv6 bge0 \  
fd4b:454e:205a:111:21e:c9ff:fed9:2f8d 64
```

- 9 Delete a previously added IPv6 virtual IP address from the configuration, enter:

```
# cfsshare deletevipv6 ipv6_address
```

For example:

```
# cfsshare deletevipv6 fd4b:454e:205a:111:21e:c9ff:fed9:2f8d
```

- 10 Unshare CFS file system, enter:

```
# cfsshare unshare mount_point
```

For example:

```
# cfsshare unshare /mnt1
```

- 11 Unmount and remove the CFS file system from the VCS configuration, enter

```
# cfsumount mount_point  
# cfsmntadm delete mount_point
```

For example:

```
# cfsumount /mnt1  
# cfsmntadm delete /mnt1
```

- 12 Unconfigure NFS shared CFS file systems, enter:

```
# cfsshare unconfig -p nfs
```


Sample 2 uses the `cfsshare` command to add a CFS file system to VCS configuration

This sample is intended to use the `cfsshare` command to add a CFS file system to VCS configuration and mount them. Then share them through NFS, unshare, unmount, and remove the CFS file systems from VCS configuration.

To configure Clustered NFS (Sample 2)

- 1 Configure a VCS configuration for CFS/CVM, enter:

```
# cfscluster config
```

- 2 Configure the CNFS components, enter:

```
# cfsshare config -p nfs shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p nfs cfsdg vollocks /locks
```

- 3 Add and mount the NFS shared CFS file system to the VCS configuration, enter:

```
# cfsshare add [-D] shared_disk_group shared_volume mount_point \  
[share_options] all=[mount_options]
```

For example:

```
# cfsshare add cfsdg vol1 /mnt1 all=rw
```

- 4 Add the virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvip [ -a nodename ] \  
network_interface address netmask
```

For example:

```
# cfsshare addvip bge0 \  
10.182.111.161 255.255.240.0
```

- 5 Delete a previously added virtual IP address from the configuration, enter:

```
# cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.111.161
```

- 6 Add the IPv6 virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvipv6 [ -a nodename ] \  
network_interface ipv6_address prefixlen
```

For example:

```
# cfsshare addvipv6 bge0 \  
fd4b:454e:205a:111:21e:c9ff:fed9:2f8d 64
```

- 7 Delete a previously added IPv6 virtual IP address from the configuration, enter:

```
# cfsshare deletevipv6 ipv6_address
```

For example:

```
# cfsshare deletevipv6 fd4b:454e:205a:111:21e:c9ff:fed9:2f8d
```

- 8 Unshare, unmount, and remove the CFS file system from the VCS configuration, enter:

```
# cfsshare delete mount_point
```

For example:

```
# cfsshare delete /mnt1
```

- 9 Unconfigure CNFS components, enter:

```
# cfsshare unconfig -p nfs
```

Sample main.cf file

This is a sample `main.cf` file.

```
include "OracleASMTypes.cf"  
include "types.cf"
```

```
include "ApplicationNone.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster cfs782 (
    UserNames = { admin = ghiAhcHeiDiiGqiChf }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
)

system system01 (
)

system system02 (
)

system system03 (
)

system system04 (
)

group cfsnfssg (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02, system03,
                     system04 }
)

Application Samba_winbind (
    StartProgram = "/opt/VRTSvcs/bin/ApplicationNone/
winbindmonitor.sh start"
    StopProgram = "/opt/VRTSvcs/bin/ApplicationNone/
winbindmonitor.sh stop"
    PidFiles = { "/var/run/winbindmonitor.pid" }
)
```

```
ApplicationNone app (
)

CFSMount cfsmount3 (
    Critical = 0
    MountPoint = "/mnt2"
    BlockDevice = "/dev/vx/dsk/fsp701-704-v03/vol2"
    NodeList = { system01, system02, system03,
                system04 }
)

CFSMount cfsnfs_locks (
    Critical = 0
    MountPoint = "/lock"
    BlockDevice = "/dev/vx/dsk/fsp701-704-v03/lock"
    NodeList = { system01, system02, system03,
                system04 }
)

CVMVolDg cvmvoldg3 (
    Critical = 0
    CVMDiskGroup = fsp701-704-v03
    CVMVolume = { lock, vol2 }
    CVMActivation @system01 = sw
    CVMActivation @system02 = sw
    CVMActivation @system03 = sw
    CVMActivation @system04 = sw
)

NFS nfs (
)

NetBios Samba_netbios (
    SambaServerRes = SambaServerResource
    NetBiosName = cfs782
)

SambaServer SambaServerResource (
    ConfFile = "/opt/pware/lib/smb.conf"
    SambaTopDir = "/opt/pware"
    LockDir = "/opt/pware/var/locks"
)
```

```
SambaShare sambashare1 (
    Critical = 0
    SambaServerRes = SambaServerResource
    ShareName = cifs1
    ShareOptions = "path=/mnt2;msdfs root=yes;msdfs
'proxy=\\10.209.116.87\\cifs1_dfs\\"
)
```

```
requires group cvm online local firm
Samba_winbind requires Samba_netbios
cfsmount3 requires cfsnfs_locks
cfsmount3 requires cvmvoldg3
cfsnfs_locks requires cvmvoldg3
sambashare1 requires SambaServerResource
sambashare1 requires cfsmount3
```

```
// resource dependency tree
//
//   group cfsnfssg
//   {
//       Application Samba_winbind
//       {
//           NetBios Samba_netbios
//       }
//       ApplicationNone app
//       NFS nfs
//       SambaShare sambashare1
//       {
//           SambaServer SambaServerResource
//           CFMount cfsmount3
//           {
//               CFMount cfsnfs_locks
//               {
//                   CVMVolDg cvmvoldg3
//               }
//               CVMVolDg cvmvoldg3
//           }
//       }
//   }
```

```
group cfsnfssg_dummy (
```

```
SystemList = { system01 = 0, system02 = 1, system03 = 2,
               system04 = 3 }
AutoFailOver = 0
Parallel = 1
AutoStartList = { system01, system02, system03,
                  system04 }
)

requires group cvm online local firm

// resource dependency tree
//
//   group cfsnfssg_dummy
//   {
//   }

group cvm (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02, system03,
                     system04 }
)

CFSfsckd vxfsckd (
    ActivationMode @system01 = { fsp701-704-v03 = sw }
    ActivationMode @system02 = { fsp701-704-v03 = sw }
    ActivationMode @system03 = { fsp701-704-v03 = sw }
    ActivationMode @system04 = { fsp701-704-v03 = sw }
)

CVMCluster cvm_clus (
    CVMClustName = cfs782
    CVMNodeId = { system01 = 0, system02 = 1,
                 system03 = 2,
                 system04 = 3 }
    CVMTransport = gab
    CVMTimeout = 200
)
```

```

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

```

```

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

```

```

// resource dependency tree
//
//   group cvm
//   {
//       CFSfsckd vxfsckd
//       {
//           CVMCluster cvm_clus
//           {
//               CVMVxconfigd cvm_vxconfigd
//           }
//       }
//   }

```

```

group vip1 (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                    system04 = 3 }
    AutoStartList = { system01, system02, system03,
                       system04 }
    TriggerPath = "bin/cavftriggers/vip"
    TriggersEnabled @system01 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system02 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system03 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system04 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    PreOnline @system01 = 1
    PreOnline @system02 = 1
    PreOnline @system03 = 1
    PreOnline @system04 = 1
)

IP vip1 (
    Device = /dev/bge:0
    Address = "10.209.116.87"
    NetMask = "255.255.252.0"
)

```

```

    )

NIC nic1 (
    Device = /dev/bge:0
    NetworkHosts = { "10.209.113.1" }
)

SambaShare sambashare1_dfs (
    Critical = 0
    SambaServerRes = SambaServerResource
    ShareName = cifs1_dfs
    ShareOptions = "path=/mnt2;readonly=no"
)

requires group cfsnfssg online local firm
sambashare1_dfs requires vip1
vip1 requires nic1

// resource dependency tree
//
//   group vip1
//   {
//     SambaShare sambashare1_dfs
//     {
//       IP vip1
//       {
//         NIC nic1
//       }
//     }
//   }
//   }

```

How to mount an NFS-exported file system on the NFS clients

This section describes how to mount an NFS-exported file system on the NFS clients.

To mount an NFS-exported file system on the NFS clients

- ◆ Run the following command:

```
# mount -F nfs -o vers=3 VIP_address:remote_filesystem mount_point
```


Disabling SMF for NFS daemons on Solaris 10

On Solaris 10, you must disable the Service Management Facility (SMF) for the NFS daemons.

UNIX start-up scripts and configuration files previously performed these functions. SMF maintains the Service Configuration Repository to store persistent configuration information as well as runtime data for all the services. All NFS daemons (`nfsd`, `mountd`, etc.) are now controlled by SMF. To keep these daemons under VCS control, modify the configuration repository to disable the SMF framework for NFS daemons.

You must invoke the following commands before bringing the NFS agent online or the agents returns an UNKNOWN state.

To keep `nfsd` and `mountd` daemons under VCS control, on all cluster nodes:

- 1 Set the `auto_enable` property to false.

```
# svccfg -s nfs/server setprop "application/auto_enable = false"
```

- 2 Refresh the SMF configuration repository.

```
# svcadm refresh nfs/server
```

- 3 Disable SMF.

```
# svcadm disable svc:/network/nfs/server:default
```

- 4 Run `svcs -a | grep -i nfs` command and review its output to make sure that SMF for `nfsd` and `mountd` is disabled.

```
# svcs -a | grep -i nfs
disabled May_29 svc:/network/nfs/server:default
```

To keep `nfsmapid` daemon under VCS control, on all cluster nodes:

- 1 Set the `auto_enable` property to false.

```
# svccfg -s nfs/mapid setprop "application/auto_enable = false"
```

- 2 Refresh the SMF configuration repository.

```
# svcadm refresh nfs/mapid
```

3 Disable SMF.

```
# svcadm disable svc:/network/nfs/mapid:default
```

4 Run `svcs -a | grep -i mapid` command and review its output to make sure that SMF for `nfsmapid` is disabled.

```
# svcs -a | grep -i mapid
disabled May_29 svc:/network/nfs/mapid:default
```

To keep the `statd` daemon under VCS control, on all cluster nodes:

1 Set the `auto_enable` property to false.

```
# svccfg -s nfs/status setprop "application/auto_enable = false"
```

2 Refresh the SMF configuration repository.

```
# svcadm refresh nfs/status
```

3 Disable SMF.

```
# svcadm disable svc:/network/nfs/status:default
```

4 Run `svcs -a | grep -i nfs` command and review its output to make sure that SMF for `statd` is disabled.

```
# svcs -a | grep -i nfs
disabled May_29 svc:/network/nfs/status:default
```

To keep `lockd` daemon under VCS control, on all cluster nodes:

1 Set the `auto_enable` property to false.

```
# svccfg -s nfs/nlockmgr setprop "application/auto_enable = false"
```

2 Refresh the SMF configuration repository.

```
# svcadm refresh nfs/nlockmgr
```

3 Disable SMF.

```
# svcadm disable svc:/network/nfs/nlockmgr:default
```

4 Run `svcs -a | grep -i nfs` command and review its output to make sure that SMF for `nlockmgr` is disabled.

```
# svcs -a | grep -i nlockmgr
disabled May_29 svc:/network/nfs/nlockmgr:default
```

Debugging Clustered NFS

The `cfsshare` command logs error messages to the VCS logs in the `/var/VRTSvcs/log` directory. The `fsclustadm frlpause_query` command may be used to display the current local copy of the global state. This may be useful in debugging any issues with this Clustered NFS feature.

Common Internet File System

- [Chapter 21. Common Internet File System](#)

Common Internet File System

This chapter includes the following topics:

- [About Common Internet File System](#)
- [Requirements for Common Internet File System](#)
- [Understanding how Samba works](#)
- [Configuring Clustered NFS and Common Internet File System on CFS](#)
- [cfsshare manual page](#)
- [Configuring Common Internet File System in user mode](#)
- [Configuring Common Internet File System in domain mode](#)
- [Configuring Common Internet File System in ads mode](#)
- [Administering Common Internet File System](#)
- [Debugging Common Internet File System](#)

About Common Internet File System

The Common Internet File System (CIFS) feature lets you share CFS file systems using CIFS protocol that can be accessed by Window clients. Upon node failure or service group failover, the CIFS shares continue to be served by other cluster nodes.

Use the `cfsshare` command to configure your CIFS shares.

See [“cfsshare manual page”](#) on page 538.

Requirements for Common Internet File System

- For Solaris 9:
 - Common Internet File System (CIFS) requires Samba version 3.2 or later and Oracle patch 112837-21.
 - If you are running Samba 3.2.x, CIFS requires you to add the Samba library path to the `LD_LIBRARY_PATH` variable in the `/opt/VRTSvcs/bin/vcsenv` file.
- For Solaris 10:
 - CIFS requires Samba version 3.2 or later.
 - If you are running Samba 3.2.x, CIFS requires you to add the Samba library path to the `LD_LIBRARY_PATH` variable in the `/opt/VRTSvcs/bin/vcsenv` file.
 - Prior knowledge of Samba is a prerequisite.

Understanding how Samba works

Samba is a networking tool that enables a UNIX server to participate in Windows networks. There are two parts to Samba, one being the server which shares out files and printers for other PC's to use, and the other being the client utilities, which allows a UNIX system to access files and printers on other Windows or Samba servers.

Configuring Clustered NFS and Common Internet File System on CFS

You can configure CNFS and CIFS on the same cluster. However, any given CFS file system can be shared using only one of the protocols (NFS or CIFS) at any given time.

cfsshare manual page

This CIFS feature adds a new functionality to the `cfsshare` utility in the VRTScavf package and several scripts that are added into the VCS configuration to manage parallel NFS and CIFS server resources. The `cfsshare` command modifies VCS resources that were created by other utilities such as `cfsmntadm`.

Note: The `cfsshare` command takes `nfs` protocol by default if the `-p` option is not used.

See the `cfsshare(1M)` manual page.

```
man -M /opt/VRTS/man cfsshare
```

Configuring Common Internet File System in user mode

This section describes how to configure CIFS in user mode.

In this mode, user authentication happens on the cluster nodes itself.

You must have NIS or some other mechanism configured on the cluster nodes to ensure the same users/groups have the same user/groups IDs on all cluster nodes.

A shared file system needs to be specified during the config operation. This file system is used to store the `smbpasswd` file, which contains the encrypted passwords for users. This way, users for whom passwords have been created on one cluster node, can authenticate themselves against other cluster nodes as well.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \  
samba_config_file -t samba_topdir -m user \  
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m user -l /var/run \  
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smbd -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = user
passwd backend = smbpasswd
smbpasswd file = pvtdir/smbpasswd
```

where *pvt*dir is the private directory of your Samba installation.

- 2 Run the following command to backup your existing `smbpasswd` file:

```
# cp -f pvtdir/smbpasswd pvtdir/smbpasswd.OLD
```

- 3 Create a symlink in *pvt*dir, that points to the `smbpasswd` file created in the locks file system:

```
# ln -sf mntpt/cifs/smbpasswd pvtdir/smbpasswd
```

where *mnt*pt is the mount point.

To unconfigure CIFS:

```
# cfsshare unconfig -p cifs
```

Note: The unconfigure operation fails if any file systems or VIPs are being shared via CIFS protocol. Use the `cfsshare delete` command to stop sharing the file systems via CIFS. Use the `cfsshare deletevip` command to remove the VIPs.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m user -l /var/run \
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```


Configuring Common Internet File System in domain mode

This section describes how to configure CIFS in domain mode.

In this mode, user authentication happens on the NT4-style Domain controllers.

The cluster nodes act as member servers in the domain. You must have performed additional steps to make user and group mapping via winbind work.

A shared file system needs to be specified during the config operation. This file system is used to replicate the `secrets.tdb` file (machine password file) across all cluster nodes. Only one of the cluster nodes joins the domain using the cluster name. Once you have copied this file to all the cluster nodes, the Domain controller sees all cluster nodes as one member server.

The shared file system can also be used to store any `tdb` files that needs to be shared across all cluster nodes. Appropriate symlinks must be created on all cluster nodes.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \  
samba_config_file -t samba_topdir -m domain \  
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m domain -l \  
/var/run -c /etc/samba/smb.conf -t /usr -s sfstest-ad \  
-d SFSTEST-AD2 -u Administrator lockdg vollocks /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smdb -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = domain
workgroup = domainname
password server = Domain_Controller_of_the_domain
```

- 2 Run the following command to backup your existing `secrets.tdb` file:

```
# mv -f pvtmdir/secrets.tdb pvtmdir/secrets.tdb.OLD
```

where `pvtmdir` is the private directory of your Samba installation.

- 3 Copy the `secrets.tdb` file created in the locks file system to the private directory of your Samba installation:

```
# cp -f mntpt/cifs/secrets.tdb pvtmdir/secrets.tdb
```

where `mntpt` is the mount point.

To unconfigure CIFS:

```
# cfsshare unconfig -p cifs
```

The unconfigure operation fails if any file systems are being shared via CIFS protocol.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m domain -l \
/var/run -c /etc/samba/smb.conf -t /usr -s sfstest-ad \
-d SFSTEST-AD2 -u Administrator lockdg vollocks /locks
```

Configuring Common Internet File System in ads mode

This section describes how to configure CIFS in ads mode.

In this mode, user authentication happens on the active directory using Kerberos. The cluster nodes act as member servers in the domain. You must have performed additional steps to make user and group mapping via winbind work.

A shared file system needs to be specified during the config operation. This file system is used to replicate the `secrets.tdb` file (machine password file) across all cluster nodes. Only one of the cluster nodes joins the domain using the cluster

name. Once you have copied this file to all the cluster nodes, the domain controller sees all cluster nodes as one member server.

You must have configured Kerberos on all cluster nodes. The time on all cluster nodes needs to be synced up with the AD server/KDC.

The shared file system can also be used to store any `tdb` file that needs to be shared across all cluster nodes. Appropriate symlinks must be created on all cluster nodes.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \
samba_config_file -t samba_topdir -m ads \
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m ads -l /var/run \
-c /etc/samba/smb.conf -t /usr -s sfstest-ad -d \
SFSTEST-AD2 -r SFSTEST-AD2.LOCAL -u Administrator \
lockdg lockvol /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smbd -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = ads
workgroup = domainname
password server = AD_server_of_the_domain
realm = realm_name
```

- 2 Run the following command to backup your existing `secrets.tdb` file:

```
# mv -f pvtmdir/secrets.tdb pvtmdir/secrets.tdb.OLD
```

where `pvtmdir` is the private directory of your Samba installation.

- 3 Copy the `secrets.tdb` file created in the locks file system to the private directory of your Samba installation:

```
# cp -f mntpt/cifs/secrets.tdb pvtmdir/secrets.tdb
```

where `mntpt` is the mount point.

To unconfigure CIFS:

```
# cfsshare unconfig -p cifs
```

Note: The unconfigure operation fails if any file systems are being shared via CIFS protocol.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m ads -l /var/run \
-c /etc/samba/smb.conf -t /usr -s sfstest-ad -d \
SFSTEST-AD2 -r SFSTEST-AD2.LOCAL -u Administrator \
lockdg lockvol /locks
```

Administering Common Internet File System

To be able to access a CIFS exported file system from a Windows client you must first add a virtual IP. You must specify this virtual IP while sharing a file system via CIFS.

Adding a virtual IP:

```
cfsshare addvip [ -a nodename ] device address netmask [networkhosts]
```

For example:

```
# cfsshare addvip bge0 10.182.79.216 \
255.255.240.0 10.182.79.215
```

The `cfsshare addvip` command lets you specify only one network interface, that is assumed to be present on all cluster nodes. If you want to specify different network interfaces for different cluster nodes, then you need to run certain VCS commands. For example:

```
# haconf -makerw
# hares -local vip1 Device
# hares -modify vip1 Device bge1 -sys system01
# hares -modify vip1 Device bge2 -sys system02
# hares -local nic1 Device
# hares -modify nic1 Device bge1 -sys system01
# hares -modify nic1 Device bge2 -sys system02
# haconf -dump -makero
```

where *vip1* is the virtual IP resource created by the `cfsshare addvip` command.

where *nic1* is the NIC resource respectively created by the `cfsshare addvip` command.

where *system01* and *system02* are the cluster nodes.

Adding and sharing a file system via CIFS:

```
cfsshare add -p cifs [-D] -v address -n cifs_share_name \
shared_disk_group shared_volume mount_point \
share_options all=[ mount_options ]
```

For example:

```
# cfsshare add -p cifs -v 10.182.79.216 \
-n sh1 sharedg vol1 /mnt1 "readonly=no" all=
```

Note: You can also use the `-c` option to specify the CIFS share options.

See the `cfsshare(1M)` manual page for more information.

Run the following command to the display the currently exported file systems:

```
# cfsshare display
```

```
CNFS metadata filesystem : /locks
Protocols Configured : NFS
```

#RESOURCE	MOUNTPOINT	PROTOCOL	OPTIONS
share1	/defragvol	NFS	rw,no_root_squash

Deleting a previous shared file system:

```
cfsshare delete -p cifs mount_point
```

For example:

```
# cfsshare delete -p cifs /mnt1
```

Deleting the VIP added previously:

```
cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.79.216
```

Sharing a CFS file system previously added to VCS

Use one of the following commands:

```
cfsshare share -p cifs -v address -n cifs share name \  
[-C cifs_share_options] mount_point
```

For example:

```
# cfsshare share -p cifs -v 10.182.79.216 -n sh1 -C readonly=no /mnt1
```

Or

```
cfsshare share -p cifs -v address -n cifs share name \  
mount_point [share_options]
```

For example:

```
# cfsshare share -p cifs -v 10.182.79.216 -n sh1 /mnt1 readonly=no
```

Note: You must add the *address* using the `cfsshare addvip` command before running the `cfsshare share` command.

See the `cfsshare(1M)` manual page for more information.

Before running this command, you should have run `cfsmntadm` command to add the shared file system to VCS configuration and the `cfsmount` command to mount the shared file system at the *mount_point*. Once these commands have been executed, the CFSMount resource corresponding to the *mount_point* gets created

in either a default service group (with a name similar to `vrts_vea_cfs_int_cfsmountnumber`) or in a separate service group, as specified by the user.

The `cfsshare share` command moves the CFSMount resource corresponding to the `mount_point` and the associated CVMVolDg resource to the `cfsnfssg` service group (that was created using the `config` option). This command also creates a share resource on top of the CFSMount resource in the same `cfsnfssg` service group.

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSMount and CVMVolDg resources in the `cfsnfssg` service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Unsharing the previous shared CFS file system

Before running this command, you should have run the `cfsshare share` command.

The `cfsshare unshare` command enables the user to stop sharing the file system mounted at the `mount_point`. This command moves the Share, CFSMount, and CVMVolDg resources corresponding to the `mount_point` from `cfsnfssg` service group to a newly created service group. The SambaShare resource is taken offline and then deleted.

```
cfsshare unshare mount_point
```

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSMount and CVMVolDg resources in the newly created service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Running the `cfsmntadm delete` command does not remove the `ActivationMode` attribute. If no volumes or vsets in the disk group are in the VCS config, you must use the `cfsdgadm delete` to remove this `ActivationMode` attribute.

Sample main.cf file

This is a sample `main.cf` file.

```
include "OracleASMTypes.cf"
include "types.cf"
include "ApplicationNone.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster cfs782 (
    UserNames = { admin = ghiAhcHeiDiiGqiChf }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
)

system system01 (
)

system system02 (
)

system system03 (
)

system system04 (
)

group cfsnfssg (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02, system03,
                     system04 }
)

Application Samba_winbind (
    StartProgram = "/opt/VRTSvcs/bin/ApplicationNone/
winbindmonitor.sh start"
    StopProgram = "/opt/VRTSvcs/bin/ApplicationNone/
winbindmonitor.sh stop"
    PidFiles = { "/var/run/winbindmonitor.pid" }
```



```
)

ApplicationNone app (
)

CFSMount cfsmount3 (
    Critical = 0
    MountPoint = "/mnt2"
    BlockDevice = "/dev/vx/dsk/fsp701-704-v03/vol2"
    NodeList = { system01, system02, system03,
                system04 }
)

CFSMount cfsnfs_locks (
    Critical = 0
    MountPoint = "/lock"
    BlockDevice = "/dev/vx/dsk/fsp701-704-v03/lock"
    NodeList = { system01, system02, system03,
                system04 }
)

CVMVolDg cvmvoldg3 (
    Critical = 0
    CVMDiskGroup = fsp701-704-v03
    CVMVolume = { lock, vol2 }
    CVMActivation @system01 = sw
    CVMActivation @system02 = sw
    CVMActivation @system03 = sw
    CVMActivation @system04 = sw
)

NFS nfs (
)

NetBios Samba_netbios (
    SambaServerRes = SambaServerResource
    NetBiosName = cfs782
)

SambaServer SambaServerResource (
    ConfFile = "/opt/pware/lib/smb.conf"
    SambaTopDir = "/opt/pware"
    LockDir = "/opt/pware/var/locks"
```

```

    )

SambaShare sambashare1 (
    Critical = 0
    SambaServerRes = SambaServerResource
    ShareName = cifs1
    ShareOptions = "path=/mnt2;msdfs root=yes;msdfs
    proxy=\\10.209.116.87\\cifs1_dfs\\"
)

requires group cvm online local firm
Samba_winbind requires Samba_netbios
cfsmount3 requires cfsnfs_locks
cfsmount3 requires cvmvoldg3
cfsnfs_locks requires cvmvoldg3
sambashare1 requires SambaServerResource
sambashare1 requires cfsmount3

// resource dependency tree
//
//   group cfsnfssg
//   {
//     Application Samba_winbind
//     {
//       NetBios Samba_netbios
//     }
//     ApplicationNone app
//     NFS nfs
//     SambaShare sambashare1
//     {
//       SambaServer SambaServerResource
//       CFSSMount cfsmount3
//       {
//         CFSSMount cfsnfs_locks
//         {
//           CVMVolDg cvmvoldg3
//         }
//         CVMVolDg cvmvoldg3
//       }
//     }
//   }

```

```
group cfsnfssg_dummy (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02, system03,
                     system04 }
)

requires group cvm online local firm

// resource dependency tree
//
//   group cfsnfssg_dummy
//   {
//   }
```

```
group cvm (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02, system03,
                     system04 }
)

CFSfsckd vxfsckd (
    ActivationMode @system01 = { fsp701-704-v03 = sw }
    ActivationMode @system02 = { fsp701-704-v03 = sw }
    ActivationMode @system03 = { fsp701-704-v03 = sw }
    ActivationMode @system04 = { fsp701-704-v03 = sw }
)

CVMCluster cvm_clus (
    CVMClustName = cfs782
    CVMNodeId = { system01 = 0, system02 = 1,
                 system03 = 2,
                 system04 = 3 }
    CVMTransport = gab
    CVMTimeout = 200
```

```

    )

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
//   group cvm
//   {
//       CFSfsckd vxfsckd
//       {
//           CVMCluster cvm_clus
//           {
//               CVMVxconfigd cvm_vxconfigd
//           }
//       }
//   }

group vip1 (
    SystemList = { system01 = 0, system02 = 1, system03 = 2,
                  system04 = 3 }
    AutoStartList = { system01, system02, system03,
                     system04 }
    TriggerPath = "bin/cavftriggers/vip"
    TriggersEnabled @system01 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system02 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system03 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    TriggersEnabled @system04 = { PREONLINE, POSTONLINE, POSTOFFLINE }
    PreOnline @system01 = 1
    PreOnline @system02 = 1
    PreOnline @system03 = 1
    PreOnline @system04 = 1
)

IP vip1 (
    Device = /dev/bge:0

```

```
        Address = "10.209.116.87"
        NetMask = "255.255.252.0"
    )

    NIC nic1 (
        Device = /dev/bge:0
        NetworkHosts = { "10.209.113.1" }
    )

    SambaShare sambashare1_dfs (
        Critical = 0
        SambaServerRes = SambaServerResource
        ShareName = cifs1_dfs
        ShareOptions = "path=/mnt2;readonly=no"
    )

    requires group cfsnfssg online local firm
    sambashare1_dfs requires vip1
    vip1 requires nic1

    // resource dependency tree
    //
    //     group vip1
    //     {
    //         SambaShare sambashare1_dfs
    //         {
    //             IP vip1
    //             {
    //                 NIC nic1
    //             }
    //         }
    //     }
    // }
```

Debugging Common Internet File System

The `cfsshare` command logs error messages to the VCS engine logs in the `/var/VRTSvcS/log` directory.

Using Point-in-time copies

- [Chapter 22. Understanding point-in-time copy methods](#)
- [Chapter 23. Administering volume snapshots](#)
- [Chapter 24. Administering Storage Checkpoints](#)
- [Chapter 25. Administering FileSnaps](#)
- [Chapter 26. Administering snapshot file systems](#)

Understanding point-in-time copy methods

This chapter includes the following topics:

- [About point-in-time copies](#)
- [When to use point-in-time copies](#)
- [About Storage Foundation point-in-time copy technologies](#)
- [Volume-level snapshots](#)
- [Storage Checkpoints](#)
- [About FileSnaps](#)
- [About snapshot file systems](#)

About point-in-time copies

Veritas Storage Foundation offers a flexible and efficient means of managing business-critical data. Storage Foundation lets you capture an online image of an actively changing database at a given instant, called a point-in-time copy.

More and more, the expectation is that the data must be continuously available (24x7) for transaction processing, decision making, intellectual property creation, and so forth. Protecting the data from loss or destruction is also increasingly important. Formerly, data was taken out of service so that the data did not change while data backups occurred; however, this option does not meet the need for minimal down time.

A point-in-time copy enables you to maximize the online availability of the data. You can perform system backup, upgrade, or perform other maintenance tasks on

the point-in-time copies. The point-in-time copies can be processed on the same host as the active data, or a different host. If required, you can offload processing of the point-in-time copies onto another host to avoid contention for system resources on your production server. This method is called off-host processing. If implemented correctly, off-host processing solutions have almost no impact on the performance of the primary production system.

When to use point-in-time copies

The following typical activities are suitable for point-in-time copy solutions implemented using Veritas FlashSnap:

- **Data backup** —Many enterprises require 24 x 7 data availability. They cannot afford the downtime involved in backing up critical data offline. By taking snapshots of your data, and backing up from these snapshots, your business-critical applications can continue to run without extended downtime or impacted performance.
- **Providing data continuity** —To provide continuity of service in the event of primary storage failure, you can use point-in-time copy solutions to recover application data. In the event of server failure, you can use point-in-time copy solutions in conjunction with the high availability cluster functionality of Veritas Storage Foundation™ Cluster File System HA or Veritas Storage Foundation HA.
- **Decision support analysis and reporting**—Operations such as decision support analysis and business reporting may not require access to real-time information. You can direct such operations to use a replica database that you have created from snapshots, rather than allow them to compete for access to the primary database. When required, you can quickly resynchronize the database copy with the data in the primary database.
- **Testing and training**—Development or service groups can use snapshots as test data for new applications. Snapshot data provides developers, system testers and QA groups with a realistic basis for testing the robustness, integrity and performance of new applications.
- **Database error recovery**—Logic errors caused by an administrator or an application program can compromise the integrity of a database. You can recover a database more quickly by restoring the database files by using Storage Checkpoints or a snapshot copy than by full restoration from tape or other backup media.
Use Storage Checkpoints to quickly roll back a database instance to an earlier point in time.
- **Cloning data**—You can clone your file system or application data. This functionality enable you to quickly and efficiently provision virtual desktops.

All of the snapshot solutions mentioned above are also available on the disaster recovery site, in conjunction with Veritas Volume Replicator.

For more information about snapshots with replication, see the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Veritas Storage Foundation provides several point-in-time copy solutions that support your needs, including the following use cases:

- Creating a replica database for decision support.
- Backing up and recovering a database with snapshots.
- Backing up and recovering an off-host cluster file system
- Backing up and recovering an online database.

Implementing point-in time copy solutions on a primary host

[Figure 22-1](#) illustrates the steps that are needed to set up the processing solution on the primary host.

Figure 22-1

Using snapshots and FastResync to implement point-in-time copy solutions on a primary host





Figure 22-2 Example point-in-time copy solution on a primary host


In this setup, it is recommended that separate paths (shown as 1 and 2) from separate controllers be configured to the disks containing the primary volumes and the snapshot volumes. This avoids contention for disk access, but the primary host's CPU, memory and I/O resources are more heavily utilized when the processing application is run.

Note: For space-optimized or unsynchronized full-sized instant snapshots, it is not possible to isolate the I/O pathways in this way. This is because such snapshots only contain the contents of changed regions from the original volume. If applications access data that remains in unchanged regions, this is read from the original volume.

Implementing off-host point-in-time copy solutions

Figure 22-3 illustrates that, by accessing snapshot volumes from a lightly loaded host (shown here as the OHP host), CPU- and I/O-intensive operations for online backup and decision support are prevented from degrading the performance of the primary host that is performing the main production activity (such as running a database).


Figure 22-3 Example implementation of an off-host point-in-time copy solution



Also, if you place the snapshot volumes on disks that are attached to host controllers other than those for the disks in the primary volumes, it is possible to avoid contending with the primary host for I/O resources. To implement this, paths 1 and 2 shown in the [Figure 22-3](#) should be connected to different controllers.

[Figure 22-4](#) shows an example of how you might achieve such connectivity using Fibre Channel technology with 4 Fibre Channel controllers in the primary host.

Figure 22-4 Example connectivity for off-host solution using redundant-loop access



This layout uses redundant-loop access to deal with the potential failure of any single component in the path between a system and a disk array.

Note: On some operating systems, controller names may differ from what is shown here.

Figure 22-5 shows how off-host processing might be implemented in a cluster by configuring one of the cluster nodes as the OHP node.

Figure 22-5 Example implementation of an off-host point-in-time copy solution using a cluster node





Figure 22-6 shows an alternative arrangement, where the OHP node could be a separate system that has a network connection to the cluster, but which is not a cluster node and is not connected to the cluster's private network.


Figure 22-6 Example implementation of an off-host point-in-time copy solution using a separate OHP host



Note: For off-host processing, the example scenarios in this document assume that a separate OHP host is dedicated to the backup or decision support role. For clusters, it may be simpler, and more efficient, to configure an OHP host that is not a member of the cluster.

[Figure 22-7](#) illustrates the steps that are needed to set up the processing solution on the primary host.

Figure 22-7 Implementing off-host processing solutions



Disk Group Split/Join is used to split off snapshot volumes into a separate disk group that is imported on the OHP host.

Note: As the snapshot volumes are to be moved into another disk group and then imported on another host, their contents must first be synchronized with the parent volumes. On reimporting the snapshot volumes, refreshing their contents from the original volume is speeded by using FastResync.

About Storage Foundation point-in-time copy technologies

This topic introduces the point-in-time copy solutions that you can implement using the Veritas FlashSnap™ technology. Veritas FlashSnap technology requires an enterprise product license.

Veritas FlashSnap offers a flexible and efficient means of managing business critical data. It allows you to capture an online image of actively changing data at a given instant: a point-in-time copy. You can perform system backup, upgrade and other maintenance tasks on point-in-time copies while providing continuous availability of your critical data. If required, you can offload processing of the point-in-time copies onto another host to avoid contention for system resources on your production server.

The following kinds of point-in-time copy solution are supported with an enterprise license:

- Volume-level solutions. There are several types of volume-level snapshots. These features are suitable for solutions where separate storage is desirable to create the snapshot. For example, lower-tier storage. Some of these techniques provided exceptional offhost processing capabilities.
See [“Volume-level snapshots”](#) on page 567.
- File system-level solutions use the Storage Checkpoint feature of Veritas File System. Storage Checkpoints are suitable for implementing solutions where storage space is critical for:
 - File systems that contain a small number of mostly large files.
 - Application workloads that change a relatively small proportion of file system data blocks (for example, web server content and some databases).
 - Applications where multiple writable copies of a file system are required for testing or versioning.

See [“Storage Checkpoints”](#) on page 571.

- File level snapshots.
 The FileSnap feature provides snapshots at the level of individual files.
 See [“About FileSnaps”](#) on page 580.

Comparison of Point-in-time copy solutions

The following table shows a side-by-side comparison of the Storage Foundation Point in time copy solutions.

Table 22-1

Solution	Granularity	Location of snapped data	Snapshot technique	Internal content	Exported content	Can be moved off-host	Availability
Instant full-sized snapshot	Volume	Separate volume	Copy on write/ Full copy	Changed regions >> Full volume	Read/Write volume	Yes, after synchronization	Immediate
Instant space optimized snapshot	Volume	Cache object (Separate cache volume)	Copy on write	Changed regions	Read/Write volume	No	Immediate
Linked plex break-off	Volume	Separate volume	Copy on write/ Full copy	Changed regions >> Full volume	Read/Write volume	Yes, after synchronization	Immediate
Plex break-off using vxsnap	Volume	Separate volume	Copy on write/ Full copy	Changed regions >> Full volume	Read/Write volume	Yes, after synchronization	Immediate
Traditional plex break-off using vxassist	Volume	Separate volume	Full copy	Full volume	Read/Write volume	Yes, after synchronization	After full synchronization
Storage Checkpoint	File system	Space within file system	Copy on write	Changed file system blocks	Read/Write file system	No	Immediate
File system snapshot	File system	Separate volume	Copy on write	Changed file system blocks	Read-only file system	No	Immediate

Table 22-1 (continued)

Solution	Granularity	Location of snapped data	Snapshot technique	Internal content	Exported content	Can be moved off-host	Availability
FileSnap	File	Space within file system	Copy on write/Lazy copy on write	Changed file system blocks	Read/Write file system	No	Immediate

Volume-level snapshots

A volume snapshot is an image of a Veritas Volume Manager (VxVM) volume at a given point in time. You can also take a snapshot of a volume set.

Volume snapshots allow you to make backup copies of your volumes online with minimal interruption to users. You can then use the backup copies to restore data that has been lost due to disk failure, software errors or human mistakes, or to create replica volumes for the purposes of report generation, application development, or testing.

Volume snapshots can also be used to implement off-host online backup.

Physically, a snapshot may be a full (complete bit-for-bit) copy of the data set, or it may contain only those elements of the data set that have been updated since snapshot creation. The latter are sometimes referred to as allocate-on-first-write snapshots, because space for data elements is added to the snapshot image only when the elements are updated (overwritten) for the first time in the original data set. Storage Foundation allocate-on-first-write snapshots are called space-optimized snapshots.

Persistent FastResync of volume snapshots

If persistent FastResync is enabled on a volume, VxVM uses a FastResync map to keep track of which blocks are updated in the volume and in the snapshot.

When snapshot volumes are reattached to their original volumes, persistent FastResync allows the snapshot data to be quickly refreshed and re-used. Persistent FastResync uses disk storage to ensure that FastResync maps survive both system and cluster crashes. If persistent FastResync is enabled on a volume in a private disk group, incremental resynchronization can take place even if the host is rebooted.

Persistent FastResync can track the association between volumes and their snapshot volumes after they are moved into different disk groups. After the disk

groups are rejoined, persistent FastResync allows the snapshot plexes to be quickly resynchronized.

Data integrity in volume snapshots

A volume snapshot captures the data that exists in a volume at a given point in time. As such, VxVM does not have any knowledge of data that is cached in memory by the overlying file system, or by applications such as databases that have files open in the file system. Snapshots are always crash consistent, that is, the snapshot can be put to use by letting the application perform its recovery. This is similar to how the application recovery occurs after a server crash. If the `fsgen` volume usage type is set on a volume that contains a mounted Veritas File System (VxFS), VxVM coordinates with VxFS to flush data that is in the cache to the volume. Therefore, these snapshots are always VxFS consistent and require no VxFS recovery while mounting.

For databases, a suitable mechanism must additionally be used to ensure the integrity of tablespace data when the volume snapshot is taken. The facility to temporarily suspend file system I/O is provided by most modern database software. The examples provided in this document illustrate how to perform this operation. For ordinary files in a file system, which may be open to a wide variety of different applications, there may be no way to ensure the complete integrity of the file data other than by shutting down the applications and temporarily unmounting the file system. In many cases, it may only be important to ensure the integrity of file data that is not in active use at the time that you take the snapshot. However, in all scenarios where application coordinate, snapshots are crash-recoverable.

Third-mirror break-off snapshots

A plex break-off snapshot uses an additional mirror to create the snapshot. Although you can create a plex break-off snapshot for a single plex volume, typically you take a snapshot of a mirrored volume. A mirrored volume has more than one plex or mirror, each of which is a copy of the data. The snapshot operation "breaks off" the plex, which becomes the snapshot volume. You can break off an existing plex or add a new plex specifically to serve as the snapshot mirror. Generally, you want to maintain redundancy for the original volume. If the original volume is a mirrored volume with two plexes, you add a third mirror for the snapshot. Hence, this type of snapshot is also known as a third-mirror snapshot.

The snapshot plex must be on a different disk from the existing plexes in the volume, within the same disk group. The disk must have enough disk space to contain the contents of the existing volume. If you have a one terabyte volume, you must have an additional one terabyte of disk space.

When you create the snapshot, the plexes are separated into two volumes. The original volume retains its original plex or plexes. The snapshot volume contains the snapshot plex. The original volume continues to take on I/O. The snapshot volume retains the data at the point of time when the snapshot was created, until you choose to perform processing on that volume.

You can make multiple snapshots, so you can have multiple copies of the original data.

Third-mirror break-off snapshots are suitable for write-intensive volumes (such as for database redo logs) where the copy-on-write mechanism of space-optimized or full-sized instant snapshots might degrade performance.

Space-optimized instant volume snapshots

Space-optimized snapshots do not contain complete physical images of the original data objects they represent. Space-optimized instant snapshots record changed regions in the original volume to a storage cache. As the original volume is written to, VxVM preserves its data in the cache before the write is committed. As the storage cache typically requires much less storage than the original volume, it is referred to as space-optimized. Space-optimized snapshots consume storage and I/O bandwidth in proportion to how much data on the original volume is updated during the life of the snapshot.

The benefits of space-optimized instant snapshots include immediate availability for use, quick refreshment, and easier configuration and administration. Because space-optimized snapshots consume less storage and I/O bandwidth than full-copy snapshots, you can take the snapshots much more frequently. This makes them well-suited for recovering from data corruption.

Space-optimized snapshots naturally tend to grow with age, as more of the data in the original objects changes, so they are inherently better-suited for shorter lifetimes.

Space-optimized snapshots cannot be taken off-host for auxiliary processing.

How space-optimized instant snapshots work


Space-optimized snapshots use a copy-on-write mechanism to make them immediately available for use when they are first created, or when their data is refreshed.

You can configure a single storage cache in a disk group that can be shared by all the volumes in that disk group. If so, the name of the cache that is declared must be the same for each volume's space-optimized snapshot. The cache is stored on disk and is persistent.

If the cache approaches full, configure VxVM to grow the cache automatically using any available free space in the disk group.

Figure 22-8 shows the instant space-optimized snapshot model.

Figure 22-8 Space-optimized instant snapshot creation and usage in a backup cycle



See “[Creating and managing space-optimized instant snapshots](#)” on page 601.

Choices for snapshot resynchronization

When a snapshot volume is reattached to its original volume within a shared disk group, there are two choices for resynchronizing the data in the volume:

- Resynchronize the snapshot from the original volume—updates the snapshot with data from the primary volume that has changed since the snapshot was taken. The snapshot is then again ready to be taken for the purposes of backup or decision support. This type of resynchronization is also known as refreshing the snapshot.
- Resynchronize the original volume from the snapshot—updates the original volume with data from the snapshot volume that has changed since the snapshot was taken. This may be necessary to restore the state of a corrupted database or file system, or to implement upgrades to production software, and is usually much quicker than using alternative approaches such as full restoration from backup media. This type of resynchronization is also known as restoring the snapshot from the copy or replica.

Disk group split/join

One or more volumes, such as snapshot volumes, can be split off into a separate disk group and deported. They are then ready for importing on another host that is dedicated to off-host processing. This host need not be a member of a cluster but it must have access to the disks on which the volumes are configured. At a later

stage, the disk group can be deported, re-imported, and joined with the original disk group, or with a different disk group.

Note: As space-optimized instant snapshots only record information about changed regions in the original volume, they cannot be moved to a different disk group. They are therefore unsuitable for the off-host processing applications that are described in this document.

The contents of full-sized instant snapshots must be fully synchronized with the unchanged regions in the original volume before such snapshots can be moved into a different disk group and deported from a host.

Storage Checkpoints

A Storage Checkpoint is a persistent image of a file system at a given instance in time. Storage Checkpoints use a copy-on-write technique to reduce I/O overhead by identifying and maintaining only those file system blocks that have changed since a previous Storage Checkpoint was taken. Storage Checkpoints have the following important features:

- Storage Checkpoints persist across system reboots and crashes.
- A Storage Checkpoint can preserve not only file system metadata and the directory hierarchy of the file system, but also user data as it existed when the Storage Checkpoint was taken.
- After creating a Storage Checkpoint of a mounted file system, you can continue to create, remove, and update files on the file system without affecting the image of the Storage Checkpoint.
- Unlike file system snapshots, Storage Checkpoints are writable.
- To minimize disk space usage, Storage Checkpoints use free space in the file system.

Storage Checkpoints and the Storage Rollback feature of Veritas Storage Foundation for Databases enable rapid recovery of databases from logical errors such as database corruption, missing files and dropped table spaces. You can mount successive Storage Checkpoints of a database to locate the error, and then roll back the database to a Storage Checkpoint before the problem occurred.

Symantec NetBackup for Oracle Advanced BLI Agent uses Storage Checkpoints to enhance the speed of backing up Oracle databases.

See the *Symantec NetBackup for Oracle Advanced BLI Agent System Administrator's Guide*.

How Storage Checkpoints differ from snapshots

Storage Checkpoints differ from Veritas File System snapshots in the following ways because they:

- Allow write operations to the Storage Checkpoint itself.
- Persist after a system reboot or failure.
- Share the same pool of free space as the file system.
- Maintain a relationship with other Storage Checkpoints by identifying changed file blocks since the last Storage Checkpoint.
- Can have multiple, read-only Storage Checkpoints that reduce I/O operations and required storage space because the most recent Storage Checkpoint is the only one that accumulates updates from the primary file system.
- Can restore the file system to its state at the time that the Storage Checkpoint was taken.

Various backup and replication solutions can take advantage of Storage Checkpoints. The ability of Storage Checkpoints to track the file system blocks that have changed since the last Storage Checkpoint facilitates backup and replication applications that only need to retrieve the changed data. Storage Checkpoints significantly minimize data movement and may promote higher availability and data integrity by increasing the frequency of backup and replication solutions.

Storage Checkpoints can be taken in environments with a large number of files, such as file servers with millions of files, with little adverse impact on performance. Because the file system does not remain frozen during Storage Checkpoint creation, applications can access the file system even while the Storage Checkpoint is taken. However, Storage Checkpoint creation may take several minutes to complete depending on the number of files in the file system.

How a Storage Checkpoint works

The Storage Checkpoint facility freezes the mounted file system (known as the primary fileset), initializes the Storage Checkpoint, and thaws the file system. Specifically, the file system is first brought to a stable state where all of its data is written to disk, and the freezing process momentarily blocks all I/O operations to the file system. A Storage Checkpoint is then created without any actual data; the Storage Checkpoint instead points to the block map of the primary fileset. The thawing process that follows restarts I/O operations to the file system.

You can create a Storage Checkpoint on a single file system or a list of file systems. A Storage Checkpoint of multiple file systems simultaneously freezes the file systems, creates a Storage Checkpoint on all of the file systems, and thaws the file systems. As a result, the Storage Checkpoints for multiple file systems have


the same creation timestamp. The Storage Checkpoint facility guarantees that multiple file system Storage Checkpoints are created on all or none of the specified file systems, unless there is a system crash while the operation is in progress.

Note: The calling application is responsible for cleaning up Storage Checkpoints after a system crash.


A Storage Checkpoint of the primary fileset initially contains only pointers to the existing data blocks in the primary fileset, and does not contain any allocated data blocks of its own.

Figure 22-9 shows the file system `/database` and its Storage Checkpoint. The Storage Checkpoint is logically identical to the primary fileset when the Storage Checkpoint is created, but it does not contain any actual data blocks.

Figure 22-9 Primary fileset and its Storage Checkpoint



In Figure 22-10, a square represents each block of the file system. This figure shows a Storage Checkpoint containing pointers to the primary fileset at the time the Storage Checkpoint is taken, as in Figure 22-9.

Figure 22-10 Initializing a Storage Checkpoint

The Storage Checkpoint presents the exact image of the file system by finding the data from the primary fileset. VxFS updates a Storage Checkpoint by using the copy-on-write technique.


See [“Copy-on-write”](#) on page 574.

Copy-on-write

In [Figure 22-11](#), the third data block in the primary fileset originally containing C is updated.

Before the data block is updated with new data, the original data is copied to the Storage Checkpoint. This is called the copy-on-write technique, which allows the Storage Checkpoint to preserve the image of the primary fileset when the Storage Checkpoint is taken.

Every update or write operation does not necessarily result in the process of copying data to the Storage Checkpoint because the old data needs to be saved only once. As blocks in the primary fileset continue to change, the Storage Checkpoint accumulates the original data blocks. In this example, subsequent updates to the third data block, now containing C', are not copied to the Storage Checkpoint because the original image of the block containing C is already saved.

Figure 22-11 Updates to the primary fileset

Storage Checkpoint visibility

With the `ckptautomnt` mount option, all Storage Checkpoints are made accessible automatically through a directory in the root directory of the file system that has the special name `.checkpoint`, which does not appear in directory listings. Inside this directory is a directory for each Storage Checkpoint in the file system. Each of these directories behave as a mount of the corresponding Storage Checkpoint, with the following exceptions:

- External applications, such as NFS, see the files as part of the original mount point. Thus, no additional NFS exports are necessary.
- Inode numbers exposed to applications can be made unique, depending on a mount option.

The Storage Checkpoints are automounted internally, but the operating system does not know about the automounting. This means that Storage Checkpoints cannot be mounted manually, and they do not appear in the list of mounted file systems. When Storage Checkpoints are created or deleted, entries in the Storage Checkpoint directory are automatically updated. If a Storage Checkpoint is removed with the `-f` option while a file in the Storage Checkpoint is still in use, the Storage Checkpoint is force unmounted, and all operations on the file fail with the EIO error.

If there is already a file or directory named `.checkpoint` in the root directory of the file system, such as a directory created with an older version of Veritas File System (VxFS) or when Storage Checkpoint visibility feature was disabled, the fake directory providing access to the Storage Checkpoints is not accessible. With this feature enabled, attempting to create a file or directory in the root directory with the name `.checkpoint` fails with the `EEXIST` error.

Storage Checkpoints and 64-bit inode numbers

The inode number of a file is the same across Storage Checkpoints. For example, if the file `file1` exists in a file system and a Storage Checkpoint is taken of that file system, running the `stat` command on `file1` in the original file system and in the Storage Checkpoint returns the same value in `st_ino`. The combination of `st_ino` and `st_dev` should uniquely identify every file in a system. This is usually not a problem because Storage Checkpoints get mounted separately, so `st_dev` is different. When accessing files in a Storage Checkpoint through the Storage Checkpoint visibility extension, `st_dev` is the same for all Storage Checkpoints as well as for the original file system. This means files can no longer be identified uniquely by `st_ino` and `st_dev`.

In general, uniquely identifying all files in a system is not necessary. However, there can be some applications that rely on unique identification to function properly. For example, a backup application might check if a file is hard-linked to another file by calling `stat` on both and checking if `st_ino` and `st_dev` are the same. If a backup application were told to back up two clones through the Storage Checkpoint visibility extension at the same time, the application can erroneously deduce that two files are the same even though the files contain different data.

By default, Veritas Storage Foundation (SF) does not make inode numbers unique. However, you can specify the `uniqueino` mount option to enable the use of unique 64-bit inode numbers. You cannot change this option during a remount.

About Storage Rollbacks

Each Storage Checkpoint is a consistent, point-in-time image of a file system, and Storage Rollback is the restore facility for these on-disk backups. Storage Rollback rolls back changed blocks contained in a Storage Checkpoint into the primary file system for faster database restoration.

Storage Checkpoints and Storage Rollback process

A Storage Checkpoint is a disk and I/O efficient snapshot technology for creating a "clone" of a currently mounted file system (the primary file system). Like a snapshot file system, a Storage Checkpoint appears as an exact image of the snapped file

system at the time the Storage Checkpoint was made. However, unlike a snapshot file system that uses separate disk space, all Storage Checkpoints share the same free space pool where the primary file system resides.

Note: A Storage Checkpoint can be mounted as read only or read-write, allowing access to the files as if it were a regular file system. A Storage Checkpoint is created using the `dbed_ckptcreate` command.

Initially, a Storage Checkpoint contains no data. The Storage Checkpoint only contains the inode list and the block map of the primary fileset. This block map points to the actual data on the primary file system. Because only the inode list and block map are required and no data is copied, creating a Storage Checkpoint takes only a few seconds and very little space.

A Storage Checkpoint initially satisfies read requests by finding the data on the primary file system, using its block map copy, and returning the data to the requesting process. When a write operation changes a data block in the primary file system, the old data is first copied to the Storage Checkpoint, and then the primary file system is updated with the new data. The Storage Checkpoint maintains the exact view of the primary file system at the time the Storage Checkpoint was taken. Subsequent writes to block *n* on the primary file system do not result in additional copies to the Storage Checkpoint because the old data only needs to be saved once. As data blocks are changed on the primary file system, the Storage Checkpoint gradually fills with the original data copied from the primary file system, and less and less of the block map in the Storage Checkpoint points back to blocks on the primary file system.

Storage Rollback restores a database, a tablespace, or datafiles on the primary file systems to the point-in-time image created during a Storage Checkpoint. Storage Rollback is accomplished by copying the "before" images from the appropriate Storage Checkpoint back to the primary file system. As with Storage Checkpoints, Storage Rollback restores at the block level, rather than at the file level. Storage Rollback is executed using the `dbed_ckptrollback` command.

For example:

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

Mountable Storage Checkpoints can be used for a wide range of application solutions including the following:

- Backups
- Investigations into data integrity
- Staging upgrades

- Database modifications
- Data replication solutions

If you mount a Storage Checkpoint as read-write, the command will not allow you to roll back to this Storage Checkpoint. This ensures that any Storage Checkpoint data that has been modified incorrectly cannot be a source of any database corruption. When a Storage Checkpoint is mounted as read-write, the `dbed_ckptmount` command creates a "shadow" Storage Checkpoint of and mounts this "shadow" Storage Checkpoint as read-write. This allows the database to still be rolled back to the original Storage Checkpoint.

For more information on mountable Storage Checkpoints:

Types of Storage Checkpoints

You can create the following types of Storage Checkpoints:

- [Data Storage Checkpoints](#)
- [Nodata Storage Checkpoints](#)
- [Removable Storage Checkpoints](#)
- [Non-mountable Storage Checkpoints](#)

Data Storage Checkpoints

A data Storage Checkpoint is a complete image of the file system at the time the Storage Checkpoint is created. This type of Storage Checkpoint contains the file system metadata and file data blocks. You can mount, access, and write to a data Storage Checkpoint just as you would to a file system. Data Storage Checkpoints are useful for backup applications that require a consistent and stable image of an active file system. Data Storage Checkpoints introduce some overhead to the system and to the application performing the write operation. For best results, limit the life of data Storage Checkpoints to minimize the impact on system resources.

See [“Showing the difference between a data and a nodata Storage Checkpoint”](#) on page 646.


Nodata Storage Checkpoints

A nodata Storage Checkpoint only contains file system metadata—no file data blocks. As the original file system changes, the nodata Storage Checkpoint records the location of every changed block. Nodata Storage Checkpoints use minimal system resources and have little impact on the performance of the file system because the data itself does not have to be copied.

In [Figure 22-12](#), the first block originally containing A is updated.

The original data is not copied to the Storage Checkpoint, but the changed block is marked in the Storage Checkpoint. The marker indicates which data has changed.

Figure 22-12 Updates to a nodata clone



See [“Showing the difference between a data and a nodata Storage Checkpoint”](#) on page 646.

Removable Storage Checkpoints

A removable Storage Checkpoint can self-destruct under certain conditions when the file system runs out of space.

See [“Storage Checkpoint space management considerations”](#) on page 654.

During user operations such as `create` or `mkdir`, if the file system runs out of space, removable Storage Checkpoints are deleted, even if the Storage Checkpoints are mounted. This ensures that applications can continue without interruptions due to lack of disk space. Non-removable Storage Checkpoints are not automatically removed under such `ENOSPC` conditions. Symantec recommends that you create only removable Storage Checkpoints. However, during certain administrative operations, such as `fsadm`, even if the file system runs out of space, removable Storage Checkpoints are not deleted.

Storage Checkpoints are created as non-removable by default. The default behavior can be changed so that VxFS creates removable Storage Checkpoints by using the `vxtunefs -D ckpt_removable=1` command. With the default set to create removable Storage Checkpoints, non-removable Storage Checkpoints can be created using `fsckptadm -R create ckpt_name mount_point` command.

See the `vxtunefs(1M)` and `fsckptadm(1M)` manual pages.

Non-mountable Storage Checkpoints

You can create Storage Checkpoints that cannot be mounted by using the `fsckptadm set nomount` command. The `nomount` option can be cleared using the `fsckptadm clear nomount` command.

Use non-mountable Storage Checkpoints as a security feature. This prevents other applications from accessing and modifying the Storage Checkpoint.

About FileSnaps

A FileSnap is an atomic space-optimized copy of a file in the same name space, stored in the same file system. Veritas File System (VxFS) supports snapshots on file system disk layout Version 8 and later.

FileSnaps provide an ability to snapshot objects that are smaller in granularity than a file system or a volume. The ability to snapshot parts of a file system name space is required for application-based or user-based management of data stored in a file system. This is useful when a file system is shared by a set of users or applications or the data is classified into different levels of importance in the same file system.

All regular file operations are supported on the FileSnap, and VxFS does not distinguish the FileSnap in any way.

Properties of FileSnaps

FileSnaps provide non-root users the ability to snapshot data that they own, without requiring administrator privileges. This enables users and applications to version, backup, and restore their data by scheduling snapshots at appropriate points of their application cycle. Restoring from a FileSnap is as simple as specifying a snapshot as the source file and the original file as the destination file as the arguments for the `vxfilesnap` command.

FileSnap creation locks the source file as read-only and locks the destination file exclusively for the duration of the operation, thus creating the snapshots atomically. The rest of the files in the file system can be accessed with no I/O pause while FileSnap creation is in progress. Read access to the source file is also uninterrupted

while the snapshot creation is in progress. This allows for true sharing of a file system by multiple users and applications in a non-intrusive fashion.

The name space relationship between source file and destination file is defined by the user-issued `vxfilesnap` command by specifying the destination file path. Veritas File System (VxFS) neither differentiates between the source file and the destination file, nor does it maintain any internal relationships between these two files. Once the snapshot is completed, the only shared property between the source file and destination file are the data blocks and block map shared by them.

The number of FileSnaps of a file is practically unlimited. The technical limit is the maximum number of files supported by the VxFS file system, which is one billion files per file set. When thousands of FileSnaps are created from the same file and each of these snapshot files is simultaneously read and written to by thousands of threads, FileSnaps scale very well due to the design that results in no contention of the shared blocks when unsharing happens due to an overwrite. The performance seen for the case of unsharing shared blocks due to an overwrite with FileSnaps is closer to that of an allocating write than that of a traditional copy-on-write.

In disk layout Version 8, to support block or extent sharing between the files, reference counts are tracked for each shared extent. VxFS processes reference count updates due to sharing and unsharing of extents in a delayed fashion. Also, an extent that is marked shared once will not go back to unshared until all the references are gone. This is to improve the FileSnap creation performance and performance of data extent unsharing. However, this in effect results in the shared block statistics for the file system to be only accurate to the point of the processing of delayed reclamation. In other words, the shared extent statistics on the file system and a file could be stale, depending on the state of the file system.

Concurrent I/O to FileSnaps

FileSnaps design and implementation ensures that concurrent reads or writes to different snapshots of the same file perform as if these were independent files. Even though the extents are shared between snapshots of the same file, the sharing has no negative impact on concurrent I/O.

Copy-on-write and FileSnaps

Veritas File System (VxFS) supports an option to do lazy copy-on-write when a region of a file referred to by a shared extent is overwritten. A typical copy-on-write implementation involves reading the old data, allocating a new block, copying or writing the old data to the new block synchronously, and writing the new data to the new block. This results in a worst case possibility of one or more allocating transactions, followed by a read, followed by a synchronous write and another write that conforms to the I/O behavior requested for the overwrite. This sequence makes

typical copy-on-write a costly operation. The VxFS lazy copy-on-write implementation does not copy the old data to the newly allocated block and hence does not have to read the old data either, as long as the new data covers the entire block. This behavior combined with delayed processing of shared extent accounting makes the lazy copy-on-write complete in times comparable to that of an allocating write. However, in the event of a server crash, when the server has not flushed the new data to the newly allocated blocks, the data seen on the overwritten region would be similar to what you would find in the case of an allocating write where the server has crashed before the data is flushed. This is not the default behavior and with the default behavior the data that you find in the overwritten region will be either the new data or the old data.

Reading from FileSnaps

For regular read requests, Veritas File System (VxFS) only caches a single copy of a data page in the page cache for a given shared data block, even though the shared data block could be accessed from any of the FileSnaps or the source file. Once the shared data page is cached, any subsequent requests via any of the FileSnaps or the source file is serviced from the page cache. This eliminates duplicate read requests to the disk, which results in lower I/O load on the array. This also reduces the page cache duplication, which results in efficient usage of system page cache with very little cache churning when thousands of FileSnaps are accessed.

Block map fragmentation and FileSnaps

The block map of the source file is shared by the snapshot file. When data is overwritten on a previously shared region, the block map of the file to which the write happens gets changed. In cases where the shared data extent of a source file is larger than the size of the overwrite request to the same region, the block map of the file that is written to becomes more fragmented.

Backup and FileSnaps

A full backup of a VxFS file system that has shared blocks may require as much space in the target as the number of total logical references to the physical blocks in the source file system. For example, if you have a 20 GB file from which one thousand FileSnaps were created, the total number of logical block references is approximately 20 TB. While the VxFS file system only requires a little over 20 GB of physical blocks to store the file and the file's one thousand snapshots, the file system requires over 20 TB of space on the backup target to back up the file system, assuming the backup target does not have deduplication support.

About snapshot file systems

A snapshot file system is an exact image of a VxFS file system, referred to as the snapped file system, that provides a mechanism for making backups. The snapshot is a consistent view of the file system “snapped” at the point in time the snapshot is made. You can select files to back up from the snapshot using a standard utility such as `cpio` or `cp`, or back up the entire file system image using the `vxdump` or `fscat` utilities.

You use the `mount` command to create a snapshot file system; the `mkfs` command is not required. A snapshot file system is always read-only. A snapshot file system exists only as long as the snapped file system is mounted, and the snapshot file system ceases to exist when unmounted. A snapped file system cannot be unmounted until all of its snapshots are unmounted. Although it is possible to have multiple snapshots of a file system made at different times, it is not possible to make a snapshot of a snapshot.

Note: A snapshot file system ceases to exist when unmounted. If mounted again, it is actually a fresh snapshot of the snapped file system. A snapshot file system must be unmounted before its dependent snapped file system can be unmounted. Neither the `fuser` command nor the `mount` command will indicate that a snapped file system cannot be unmounted because a snapshot of it exists.

On cluster file systems, snapshots can be created on any node in the cluster, and backup operations can be performed from that node. The snapshot of a cluster file system is accessible only on the node where it is created, that is, the snapshot file system itself cannot be cluster mounted.

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

How a snapshot file system works

A snapshot file system is created by mounting an empty disk slice as a snapshot of a currently mounted file system. The bitmap, blockmap and super-block are initialized and then the currently mounted file system is frozen. After the file system to be snapped is frozen, the snapshot is enabled and mounted and the snapped file system is thawed. The snapshot appears as an exact image of the snapped file system at the time the snapshot was made.

See [“Freezing and thawing a file system”](#) on page 489.

Initially, the snapshot file system satisfies read requests by finding the data on the snapped file system and returning it to the requesting process. When an inode

update or a write changes the data in block *n* of the snapped file system, the old data is first read and copied to the snapshot before the snapped file system is updated. The bitmap entry for block *n* is changed from 0 to 1, indicating that the data for block *n* can be found on the snapshot file system. The blockmap entry for block *n* is changed from 0 to the block number on the snapshot file system containing the old data.

A subsequent read request for block *n* on the snapshot file system will be satisfied by checking the bitmap entry for block *n* and reading the data from the indicated block on the snapshot file system, instead of from block *n* on the snapped file system. This technique is called copy-on-write. Subsequent writes to block *n* on the snapped file system do not result in additional copies to the snapshot file system, since the old data only needs to be saved once.

All updates to the snapped file system for inodes, directories, data in files, extent maps, and so forth, are handled in this fashion so that the snapshot can present a consistent view of all file system structures on the snapped file system for the time when the snapshot was created. As data blocks are changed on the snapped file system, the snapshot gradually fills with data copied from the snapped file system.

The amount of disk space required for the snapshot depends on the rate of change of the snapped file system and the amount of time the snapshot is maintained. In the worst case, the snapped file system is completely full and every file is removed and rewritten. The snapshot file system would need enough blocks to hold a copy of every block on the snapped file system, plus additional blocks for the data structures that make up the snapshot file system. This is approximately 101 percent of the size of the snapped file system. Normally, most file systems do not undergo changes at this extreme rate. During periods of low activity, the snapshot should only require two to six percent of the blocks of the snapped file system. During periods of high activity, the snapshot might require 15 percent of the blocks of the snapped file system. These percentages tend to be lower for larger file systems and higher for smaller ones.

Warning: If a snapshot file system runs out of space for changed data blocks, it is disabled and all further attempts to access it fails. This does not affect the snapped file system.

Administering volume snapshots

This chapter includes the following topics:

- [About volume snapshots](#)
- [How traditional third-mirror break-off snapshots work](#)
- [How full-sized instant snapshots work](#)
- [Linked break-off snapshot volumes](#)
- [Cascaded snapshots](#)
- [Creating multiple snapshots](#)
- [Restoring the original volume from a snapshot](#)
- [Creating instant snapshots](#)
- [Creating traditional third-mirror break-off snapshots](#)
- [Adding a version 0 DCO and DCO volume](#)

About volume snapshots

VxVM can take an image of a volume at a given point in time. This image is called a volume snapshot.

See [“Volume-level snapshots”](#) on page 567.

You can also take a snapshot of a volume set.

Snapshot creation using the `vxsnap` command is the preferred mechanism for implementing point-in-time copy solutions in VxVM. Support for traditional third-mirror

snapshots that are created using the `vxassist` command may be removed in a future release.

To recover from the failure of instant snapshot commands, see the *Veritas Storage Foundation and High Availability Troubleshooting Guide*.


How traditional third-mirror break-off snapshots work

The recommended approach to performing volume backup from the command line, or from a script, is to use the `vxsnap` command. The `vxassist snapstart`, `snapshotwait`, and `snapshot` commands are supported for backward compatibility.

The use of the `vxassist` command to administer traditional (third-mirror break-off) snapshots is not supported for volumes that are prepared for instant snapshot creation. Use the `vxsnap` command instead.

Figure 23-1 shows the traditional third-mirror break-off volume snapshot model that is supported by the `vxassist` command.

Figure 23-1 Third-mirror snapshot creation and usage



The `vxassist snapstart` command creates a mirror to be used for the snapshot, and attaches it to the volume as a snapshot mirror. As is usual when creating a mirror, the process of copying the volume's contents to the new snapshot plexes can take some time to complete. (The `vxassist snapabort` command cancels this operation and removes the snapshot mirror.)

When the attachment is complete, the `vxassist snapshot` command is used to create a new snapshot volume by taking one or more snapshot mirrors to use as

its data plexes. The snapshot volume contains a copy of the original volume's data at the time that you took the snapshot. If more than one snapshot mirror is used, the snapshot volume is itself mirrored.

The command, `vxassist snapback`, can be used to return snapshot plexes to the original volume from which they were snapped, and to resynchronize the data in the snapshot mirrors from the data in the original volume. This enables you to refresh the data in a snapshot after you use it to make a backup. You can use a variation of the same command to restore the contents of the original volume from a snapshot previously taken.

The FastResync feature minimizes the time and I/O needed to resynchronize the data in the snapshot. If FastResync is not enabled, a full resynchronization of the data is required.


Finally, you can use the `vxassist snapclear` command to break the association between the original volume and the snapshot volume. Because the snapshot relationship is broken, no change tracking occurs. Use this command if you do not need to reuse the snapshot volume to create a new point-in-time copy.

How full-sized instant snapshots work

Full-sized instant snapshots are a variation on the third-mirror volume snapshot model that make a snapshot volume available for I/O access as soon as the snapshot plexes have been created.

Figure 23-2 shows the full-sized instant volume snapshot model.

Figure 23-2 Full-sized instant snapshot creation and usage in a backup cycle



To create an instant snapshot, use the `vxsnap make` command. This command can either be applied to a suitably prepared empty volume that is to be used as the

snapshot volume, or it can be used to break off one or more synchronized plexes from the original volume.

You can make a backup of a full-sized instant snapshot, instantly refresh its contents from the original volume, or attach its plexes to the original volume, without completely synchronizing the snapshot plexes from the original volume.

VxVM uses a copy-on-write mechanism to ensure that the snapshot volume preserves the contents of the original volume at the time that the snapshot is taken. Any time that the original contents of the volume are about to be overwritten, the original data in the volume is moved to the snapshot volume before the write proceeds. As time goes by, and the contents of the volume are updated, its original contents are gradually relocated to the snapshot volume.

If a read request comes to the snapshot volume, yet the data resides on the original volume (because it has not yet been changed), VxVM automatically and transparently reads the data from the original volume.

If desired, you can perform either a background (non-blocking) or foreground (blocking) synchronization of the snapshot volume. This is useful if you intend to move the snapshot volume into a separate disk group for off-host processing, or you want to turn the snapshot volume into an independent volume.

The `vxsnap refresh` command allows you to update the data in a snapshot, for example, before taking a backup.

The command `vxsnap reattach` attaches snapshot plexes to the original volume, and resynchronizes the data in these plexes from the original volume. Alternatively, you can use the `vxsnap restore` command to restore the contents of the original volume from a snapshot that you took at an earlier point in time. You can also choose whether or not to keep the snapshot volume after restoration of the original volume is complete.

By default, the FastResync feature of VxVM is used to minimize the time and I/O needed to resynchronize the data in the snapshot mirror. FastResync must be enabled to create instant snapshots.

See [“Creating and managing full-sized instant snapshots”](#) on page 604.

An empty volume must be prepared for use by full-sized instant snapshots and linked break-off snapshots.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

Linked break-off snapshot volumes

A variant of third-mirror break-off snapshots are linked break-off snapshots, which use the `vxsnap addmir` command to link a specially prepared volume with the data volume. The volume that is used for the snapshot is prepared in the same way as for full-sized instant snapshots. However, unlike full-sized instant snapshots, this volume can be set up in a different disk group from the data volume. This makes linked break-off snapshots especially suitable for recurring off-host processing applications as it avoids the disk group split/join administrative step. As with third-mirror break-off snapshots, you must wait for the contents of the snapshot volume to be synchronized with the data volume before you can use the `vxsnap make` command to take the snapshot.

When a link is created between a volume and the mirror that will become the snapshot, separate link objects (similar to snap objects) are associated with the volume and with its mirror. The link object for the original volume points to the mirror volume, and the link object for the mirror volume points to the original volume. All I/O is directed to both the original volume and its mirror, and a synchronization of the mirror from the data in the original volume is started.

You can use the `vxprint` command to display the state of link objects, which appear as type `ln`. Link objects can have the following states:

ACTIVE	The mirror volume has been fully synchronized from the original volume. The <code>vxsnap make</code> command can be run to create a snapshot.
ATTACHING	Synchronization of the mirror volume is in progress. The <code>vxsnap make</code> command cannot be used to create a snapshot until the state changes to ACTIVE. The <code>vxsnap snapwait</code> command can be used to wait for the synchronization to complete.
BROKEN	The mirror volume has been detached from the original volume because of an I/O error or an unsuccessful attempt to grow the mirror volume. The <code>vxrecover</code> command can be used to recover the mirror volume in the same way as for a DISABLED volume.

If you resize (grow or shrink) a volume, all its `ACTIVE` linked mirror volumes are also resized at the same time. The volume and its mirrors can be in the same disk group or in different disk groups. If the operation is successful, the volume and its mirrors will have the same size.

If a volume has been grown, a resynchronization of the grown regions in its linked mirror volumes is started, and the links remain in the `ATTACHING` state until resynchronization is complete. The `vxsnap snapwait` command can be used to wait for the state to become `ACTIVE`.

When you use the `vxsnap make` command to create the snapshot volume, this removes the link, and establishes a snapshot relationship between the snapshot volume and the original volume.

The `vxsnap reattach` operation re-establishes the link relationship between the two volumes, and starts a resynchronization of the mirror volume.

See [“Creating and managing linked break-off snapshot volumes”](#) on page 609.


An empty volume must be prepared for use by linked break-off snapshots.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

Cascaded snapshots

[Figure 23-3](#) shows a snapshot hierarchy, known as a snapshot cascade, that can improve write performance for some applications.

Figure 23-3 Snapshot cascade



Instead of having several independent snapshots of the volume, it is more efficient to make the older snapshots into children of the latest snapshot.

A snapshot cascade is most likely to be used for regular online backup of a volume where space-optimized snapshots are written to disk but not to tape.

A snapshot cascade improves write performance over the alternative of several independent snapshots, and also requires less disk space if the snapshots are space-optimized. Only the latest snapshot needs to be updated when the original volume is updated. If and when required, the older snapshots can obtain the changed data from the most recent snapshot.

A snapshot may be added to a cascade by specifying the `infrontof` attribute to the `vxsnap make` command when the second and subsequent snapshots in the cascade are created. Changes to blocks in the original volume are only written to the most recently created snapshot volume in the cascade. If an attempt is made to read data from an older snapshot that does not exist in that snapshot, it is obtained by searching recursively up the hierarchy of more recent snapshots.

The following points determine whether it is appropriate to use a snapshot cascade:

- Deletion of a snapshot in the cascade takes time to copy the snapshot's data to the next snapshot in the cascade.
- The reliability of a snapshot in the cascade depends on all the newer snapshots in the chain. Thus the oldest snapshot in the cascade is the most vulnerable.
- Reading from a snapshot in the cascade may require data to be fetched from one or more other snapshots in the cascade.


For these reasons, it is recommended that you do not attempt to use a snapshot cascade with applications that need to remove or split snapshots from the cascade. In such cases, it may be more appropriate to create a snapshot of a snapshot as described in the following section.

See [“Adding a snapshot to a cascaded snapshot hierarchy”](#) on page 615.

Note: Only unsynchronized full-sized or space-optimized instant snapshots are usually cascaded. It is of little utility to create cascaded snapshots if the `infrontof` snapshot volume is fully synchronized (as, for example, with break-off type snapshots).

Creating a snapshot of a snapshot

Figure 23-4 creation of a snapshot of an existing snapshot.



Even though the arrangement of the snapshots in this figure appears similar to a snapshot cascade, the relationship between the snapshots is not recursive. When reading from the snapshot `s2`, data is obtained directly from the original volume, `v`, if it does not exist in `s1` itself.

See [Figure 23-3](#) on page 590.

Such an arrangement may be useful if the snapshot volume, `s1`, is critical to the operation. For example, `s1` could be used as a stable copy of the original volume, `v`. The additional snapshot volume, `s2`, can be used to restore the original volume if that volume becomes corrupted. For a database, you might need to replay a redo log on `s2` before you could use it to restore `v`.


Figure 23-5 shows the sequence of steps that would be required to restore a database.

Figure 23-5 Using a snapshot of a snapshot to restore a database


- 1 Create instant snapshot S1 of volume V




- 2 Create instant snapshot S2 of S1



- 3 After contents of V have gone bad, apply the database to redo logs to S2



- 4 Restore contents of V instantly from snapshot S2 and keep S1 as a stable copy




If you have configured snapshots in this way, you may wish to make one or more of the snapshots into independent volumes. There are two `vxsnap` commands that you can use to do this:

- `vxsnap dis` dissociates a snapshot and turns it into an independent volume. The snapshot to be dissociated must have been fully synchronized from its parent. If a snapshot volume has a child snapshot volume, the child must also have been fully synchronized. If the command succeeds, the child snapshot becomes a snapshot of the original volume.


Figure 23-6 shows the effect of applying the `vxsnap dis` command to snapshots with and without dependent snapshots.

Figure 23-6 Dissociating a snapshot volume

`vxsnap dis` is applied to snapshot S2, which has no snapshots of its own




`vxsnap dis` is applied to snapshot S1, which has one snapshot S2



- `vxsnap split` dissociates a snapshot and its dependent snapshots from its parent volume. The snapshot that is to be split must have been fully synchronized from its parent volume.

Figure 23-7 shows the operation of the `vxsnap split` command.

Figure 23-7 Splitting snapshots



Creating multiple snapshots

To make it easier to create snapshots of several volumes at the same time, both the `vxsnap make` and `vxassist snapshot` commands accept more than one volume name as their argument.

For traditional snapshots, you can create snapshots of all the volumes in a single disk group by specifying the option `-o allvols` to the `vxassist snapshot` command.

By default, each replica volume is named `SNAPnumber-volume`, where `number` is a unique serial number, and `volume` is the name of the volume for which a snapshot is being taken. This default can be overridden by using the option `-o name=pattern`.

See the `vxassist(1M)` manual page.

See the `vxsnap(1M)` manual page.

You can create a snapshot of all volumes that form a logical group; for example, all the volumes that conform to a database instance.


Restoring the original volume from a snapshot

For traditional snapshots, the snapshot plex is resynchronized from the data in the original volume during a `vxassist snapback` operation.

Figure 23-8 shows an alternative where the snapshot overwrites the original volume.

Figure 23-8 Resynchronizing an original volume from a snapshot

Refresh on snapback



Specifying the option `-o resyncfromreplica` to `vxassist resynchronizes` the original volume from the data in the snapshot.

Warning: The original volume must not be in use during a `snapback` operation that specifies the option `-o resyncfromreplica` to resynchronize the volume from a snapshot. Stop any application, such as a database, and unmount any file systems that are configured to use the volume.

For instant snapshots, the `vxsnap restore` command may be used to restore the contents of the original volume from an instant snapshot or from a volume derived from an instant snapshot. The volume that is used to restore the original volume can either be a true backup of the contents of the original volume at some point in time, or it may have been modified in some way (for example, by applying a database

log replay or by running a file system checking utility such as `fsck`). All synchronization of the contents of this backup must have been completed before the original volume can be restored from it. The original volume is immediately available for use while its contents are being restored.

See [“Restoring a volume from an instant space-optimized snapshot”](#) on page 618.

Creating instant snapshots

Note: You need a Storage Foundation Enterprise license to use this feature.

VxVM allows you to make instant snapshots by using the `vxsnap` command.

You can also take instant snapshots of RAID-5 volumes that have been converted to a special layered volume layout by the addition of a DCO and DCO volume.

A plex in a full-sized instant snapshot requires as much space as the original volume. If you instead make a space-optimized instant snapshot of a volume, this only requires enough storage to record the original contents of the parent volume as they are changed during the life of the snapshot.

The recommended approach to performing volume backup from the command line, or from a script, is to use the `vxsnap` command. The `vxsnap prepare` and `make` tasks allow you to back up volumes online with minimal disruption to users.

`vxsnap prepare` creates a DCO and DCO volume and associates this with the original volume. It also enables Persistent FastResync.

`vxsnap make` creates an instant snapshot that is immediately available for making a backup. After the snapshot has been taken, read requests for data in the instant snapshot volume are satisfied by reading either from a non-updated region of the original volume, or from the copy of the original contents of an updated region that have been recorded by the snapshot.

Note: Synchronization of a full-sized instant snapshot from the original volume is enabled by default. If you specify the `syncing=no` attribute to `vxsnap make`, this disables synchronization, and the contents of the instant snapshot are unlikely ever to become fully synchronized with the contents of the original volume at the point in time that the snapshot was taken. In such a case, the snapshot cannot be used for off-host processing, nor can it become an independent volume.

You can immediately retake a full-sized or space-optimized instant snapshot at any time by using the `vxsnap refresh` command. If a fully synchronized instant snapshot is required, the new resynchronization must first complete.

To create instant snapshots of volume sets, use volume set names in place of volume names in the `vxsnap` command.

See [“Creating instant snapshots of volume sets”](#) on page 612.

When using the `vxsnap prepare` or `vxassist make` commands to make a volume ready for instant snapshot operations, if the specified region size exceeds half the value of the tunable `voliomem_maxpool_sz`, the operation succeeds but gives a warning such as the following (for a system where `voliomem_maxpool_sz` is set to 12MB):

```
VxVM vxassist WARNING V-5-1-0 Specified regionsize is
larger than the limit on the system
(voliomem_maxpool_sz/2=6144k).
```

If this message is displayed, `vxsnap make`, `refresh` and `restore` operations on such volumes fail as they might potentially hang the system. Such volumes can be used only for break-off snapshot operations using the `reattach` and `make` operations.

To make the volumes usable for instant snapshot operations, use `vxsnap unprepare` on the volume, and then use `vxsnap prepare` to re-prepare the volume with a region size that is less than half the size of `voliomem_maxpool_sz` (in this example, 1MB):

```
# vxsnap -g mydg -f unprepare vol1
# vxsnap -g mydg prepare vol1 regionsize=1M
```

See [“Creating instant snapshots of volume sets”](#) on page 612.

See [“Creating and managing space-optimized instant snapshots”](#) on page 601.

See [“Creating and managing full-sized instant snapshots”](#) on page 604.

See [“Creating and managing third-mirror break-off snapshots”](#) on page 606.

See [“Creating and managing linked break-off snapshot volumes”](#) on page 609.

Adding an instant snap DCO and DCO volume

To prepare a volume for instant snapshots, an instant snap Data Change Object (DCO) and DCO volume must be associated with that volume. This procedure also enables Persistent FastResync on the volume.

The following procedure is required only if the volume does not have an instant snap DCO volume.

By default, volumes on thin provisioning LUNs are created with an instant snap DCO volume.

To an an instant snap DCO and DCO volume

- 1 Verify that the volume has an instant snap data change object (DCO) and DCO volume, and that FastResync is enabled on the volume:

```
# vxprint -g volumedg -F%instant volume
# vxprint -g volumedg -F%fastresync volume
```

If both commands return a value of `on`, skip to step 3. Otherwise continue with step 2.

- 2 To prepare a volume for instant snapshots, use the following command:

```
# vxsnap [-g diskgroup] prepare volume [regionsize=size] \
    [ndcomirs=number] [alloc=storage_attributes]
```

Run the `vxsnap prepare` command on a volume only if it does not have an instant snap DCO volume

For example, to prepare the volume, `myvol`, in the disk group, `mydg`, use the following command:

```
# vxsnap -g mydg prepare myvol regionsize=128k ndcomirs=2 \
    alloc=mydg10,mydg11
```

This example creates a DCO object and redundant DCO volume with two plexes located on disks `mydg10` and `mydg11`, and associates them with `myvol`. The region size is also increased to 128KB from the default size of 64KB. The region size must be a power of 2, and be greater than or equal to 16KB. A smaller value requires more disk space for the change maps, but the finer granularity provides faster resynchronization.

- 3 If you need several space-optimized instant snapshots for the volumes in a disk group, you may find it convenient to create a single shared cache object in the disk group rather than a separate cache object for each snapshot.

See [“Creating a shared cache object”](#) on page 598.

For full-sized instant snapshots and linked break-off snapshots, you must prepare a volume that is to be used as the snapshot volume. This volume must be the same size as the data volume for which the snapshot is being created, and it must also have the same region size.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

Creating a shared cache object

To create a shared cache object

- 1 Decide on the following characteristics that you want to allocate to the cache volume that underlies the cache object:
 - The cache volume size should be sufficient to record changes to the parent volumes during the interval between snapshot refreshes. A suggested value is 10% of the total size of the parent volumes for a refresh interval of 24 hours.
 - The cache volume can be mirrored for redundancy.
 - If the cache volume is mirrored, space is required on at least as many disks as it has mirrors. These disks should not be shared with the disks used for the parent volumes. The disks should not be shared with disks used by critical volumes to avoid impacting I/O performance for critical volumes, or hindering disk group split and join operations.
- 2 Having decided on its characteristics, use the `vxassist` command to create the cache volume. The following example creates a mirrored cache volume, `cachevol`, with size 1GB in the disk group, `mydg`, on the disks `mydg16` and `mydg17`:

```
# vxassist -g mydg make cachevol 1g layout=mirror \  
init=active mydg16 mydg17
```

The attribute `init=active` makes the cache volume immediately available for use.

- 3 Use the `vxmake cache` command to create a cache object on top of the cache volume that you created in the previous step:

```
# vxmake [-g diskgroup] cache cache_object \  
  cachevolname=volume [regionsize=size] [autogrow=on] \  
  [highwatermark=hwmk] [autogrowby=agbvalue] \  
  [maxautogrow=maxagbvalue]
```

If the region size, `regionsize`, is specified, it must be a power of 2, and be greater than or equal to 16KB (16k). If not specified, the region size of the cache is set to 64KB.

All space-optimized snapshots that share the cache must have a region size that is equal to or an integer multiple of the region size set on the cache. Snapshot creation also fails if the original volume's region size is smaller than the cache's region size.

If the region size of a space-optimized snapshot differs from the region size of the cache, this can degrade the system's performance compared to the case where the region sizes are the same.

To prevent the cache from growing automatically, specify `autogrow=off`. By default, `autogrow=on`.

In the following example, the cache object, `cobjmydg`, is created over the cache volume, `cachevol`, the region size of the cache is set to 32KB, and the `autogrow` feature is enabled:

```
# vxmake -g mydg cache cobjmydg cachevolname=cachevol \  
  regionsize=32k autogrow=on
```

- 4 Enable the cache object using the following command:

```
# vxcache [-g diskgroup] start cache_object
```

For example to start the cache object, `cobjmydg`:

```
# vxcache -g mydg start cobjmydg
```

See [“Removing a cache”](#) on page 626.

Creating a volume for use as a full-sized instant or linked break-off snapshot

To create an empty volume for use by a full-sized instant snapshot or a linked break-off snapshot

- 1 Use the `vxprint` command on the original volume to find the required size for the snapshot volume.

```
# LEN=`vxprint [-g diskgroup] -F%len volume`
```

The command as shown assumes a Bourne-type shell such as `sh`, `ksh` or `bash`. You may need to modify the command for other shells such as `csh` or `tcsh`.

- 2 Use the `vxprint` command on the original volume to discover the name of its DCO:

```
# DCONAME=`vxprint [-g diskgroup] -F%dco_name volume`
```

- 3 Use the `vxprint` command on the DCO to discover its region size (in blocks):

```
# RSZ=`vxprint [-g diskgroup] -F%regionsz $DCONAME`
```

- 4 Use the `vxassist` command to create a volume, *snapvol*, of the required size and redundancy, together with an instant snap DCO volume with the correct region size:

```
# vxassist [-g diskgroup] make snapvol $LEN \
  [layout=mirror nmirror=number] logtype=dco dnl=off \
  dcoversion=20 [ndcomirror=number] regionsz=$RSZ \
  init=active [storage_attributes]
```

Storage attributes give you control over the devices, including disks and controllers, which `vxassist` uses to configure a volume.

See [“Creating a volume on specific disks”](#) on page 204.

Specify the same number of DCO mirrors (`ndcomirror`) as the number of mirrors in the volume (`nmirror`). The `init=active` attribute makes the volume available immediately. You can use storage attributes to specify which disks should be used for the volume.

As an alternative to creating the snapshot volume and its DCO volume in a single step, you can first create the volume, and then prepare it for instant snapshot operations as shown here:

```
# vxassist [-g diskgroup] make snapvol $LEN \
  [layout=mirror nmirror=number] init=active \
  [storage_attributes]
# vxsnap [-g diskgroup] prepare snapvol [ndcomirs=number] \
  regionsize=$RSZ [storage_attributes]
```

Creating and managing space-optimized instant snapshots

Space-optimized instant snapshots are not suitable for write-intensive volumes (such as for database redo logs) because the copy-on-write mechanism may degrade performance.

To split the volume and snapshot into separate disk groups (for example, to perform off-host processing), you must use a fully synchronized full-sized instant, third-mirror break-off or linked break-off snapshot (which do not require a cache object). You cannot use a space-optimized instant snapshot.

Creation of space-optimized snapshots that use a shared cache fails if the region size specified for the volume is smaller than the region size set on the cache.

If the region size of a space-optimized snapshot differs from the region size of the cache, this can degrade the system's performance compared to the case where the region sizes are the same.

See [“Creating a shared cache object”](#) on page 598.

The attributes for a snapshot are specified as a tuple to the `vxsnap make` command. This command accepts multiple tuples. One tuple is required for each snapshot that is being created. Each element of a tuple is separated from the next by a slash character (/). Tuples are separated by white space.

To create and manage a space-optimized instant snapshot

- 1 Use the `vxsnap make` command to create a space-optimized instant snapshot. This snapshot can be created by using an existing cache object in the disk group, or a new cache object can be created.
 - To create a space-optimized instant snapshot, `snapvol`, that uses a named shared cache object:

```
# vxsnap [-g diskgroup] make source=vol/newvol=snapvol\  
/cache=cacheobject [alloc=storage_attributes]
```

For example, to create the space-optimized instant snapshot, `snap3myvol`, of the volume, `myvol`, in the disk group, `mydg`, on the disk `mydg14`, and which uses the shared cache object, `cobjmydg`, use the following command:

```
# vxsnap -g mydg make source=myvol/newvol=snap3myvol\  
/cache=cobjmydg alloc=mydg14
```

The DCO is created on the specified allocation.

- To create a space-optimized instant snapshot, `snapvol`, and also create a cache object for it to use:

```
# vxsnap [-g diskgroup] make source=vol/newvol=snapvol\  
[/cachesize=size] [/autogrow=yes] [/ncachemirror=number]\  
[alloc=storage_attributes]
```

The `cachesize` attribute determines the size of the cache relative to the size of the volume. The `autogrow` attribute determines whether VxVM will automatically enlarge the cache if it is in danger of overflowing. By default, `autogrow=on` and the cache is automatically grown.

If `autogrow` is enabled, but the cache cannot be grown, VxVM disables the oldest and largest snapshot that is using the same cache, and releases its cache space for use.

The `ncachemirror` attribute specifies the number of mirrors to create in the cache volume. For backup purposes, the default value of 1 should be sufficient.

For example, to create the space-optimized instant snapshot, `snap4myvol`, of the volume, `myvol`, in the disk group, `mydg`, on the disk `mydg15`, and which uses a newly allocated cache object that is 1GB in size, but which can automatically grow in size, use the following command:

```
# vxsnap -g mydg make source=myvol/new=snap4myvol\  
/cachesize=1g/autogrow=yes alloc=mydg15
```

If a cache is created implicitly by specifying `cachesize`, and `ncachemirror` is specified to be greater than 1, a DCO is attached to the cache volume to enable dirty region logging (DRL). DRL allows fast recovery of the cache backing store after a system crash. The DCO is allocated on the same disks as those that are occupied by the DCO of the source volume. This is done to allow the cache and the source volume to remain in the same disk group for disk group move, split and join operations.

- 2 Use `fsck` (or some utility appropriate for the application running on the volume) to clean the temporary volume's contents. For example, you can use this command with a VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/snapshot
```

The specified device must have a valid entry in the `/etc/vfstab` file.

- 3 To backup the data in the snapshot, use an appropriate utility or operating system command to copy the contents of the snapshot to tape, or to some other backup medium.
- 4 You now have the following options:
 - Refresh the contents of the snapshot. This creates a new point-in-time image of the original volume ready for another backup. If synchronization was already in progress on the snapshot, this operation may result in large portions of the snapshot having to be resynchronized.
See [“Refreshing an instant space-optimized snapshot”](#) on page 616.
 - Restore the contents of the original volume from the snapshot volume. The space-optimized instant snapshot remains intact at the end of the operation.
See [“Restoring a volume from an instant space-optimized snapshot”](#) on page 618.
 - Destroy the snapshot.
See [“Removing an instant snapshot”](#) on page 619.

Creating and managing full-sized instant snapshots

Full-sized instant snapshots are not suitable for write-intensive volumes (such as for database redo logs) because the copy-on-write mechanism may degrade the performance of the volume.

For full-sized instant snapshots, you must prepare a volume that is to be used as the snapshot volume. This must be the same size as the volume for which the snapshot is being created, and it must also have the same region size.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

The attributes for a snapshot are specified as a tuple to the `vxsnap make` command. This command accepts multiple tuples. One tuple is required for each snapshot that is being created. Each element of a tuple is separated from the next by a slash character (/). Tuples are separated by white space.

To create and manage a full-sized instant snapshot

- 1 To create a full-sized instant snapshot, use the following form of the `vxsnap make` command:

```
# vxsnap [-g diskgroup] make source=volume/snapvol=snapvol\  
[/snapdg=snapdiskgroup] [/syncing=off]
```

The command specifies the volume, *snapvol*, that you prepared earlier.

For example, to use the prepared volume, *snap1myvol*, as the snapshot for the volume, *myvol*, in the disk group, *mydg*, use the following command:

```
# vxsnap -g mydg make source=myvol/snapvol=snap1myvol
```

For full-sized instant snapshots that are created from an empty volume, background synchronization is enabled by default (equivalent to specifying the `syncing=on` attribute). To move a snapshot into a separate disk group, or to turn it into an independent volume, you must wait for its contents to be synchronized with those of its parent volume.

You can use the `vxsnap syncwait` command to wait for the synchronization of the snapshot volume to be completed, as shown here:

```
# vxsnap [-g diskgroup] syncwait snapvol
```

For example, you would use the following command to wait for synchronization to finish on the snapshot volume, *snap2myvol*:

```
# vxsnap -g mydg syncwait snap2myvol
```


This command exits (with a return code of zero) when synchronization of the snapshot volume is complete. The snapshot volume may then be moved to another disk group or turned into an independent volume.

See [“Controlling instant snapshot synchronization”](#) on page 622.

If required, you can use the following command to test if the synchronization of a volume is complete:

```
# vxprint [-g diskgroup] -F%incomplete snapvol
```

This command returns the value `off` if synchronization of the volume, *snapvol*, is complete; otherwise, it returns the value `on`.

You can also use the `vxsnap print` command to check on the progress of synchronization.

See [“Displaying snapshot information”](#) on page 634.

If you do not want to move the snapshot into a separate disk group, or to turn it into an independent volume, specify the `syncing=off` attribute. This avoids unnecessary system overhead. For example, to turn off synchronization when creating the snapshot of the volume, *myvol*, you would use the following form of the `vxsnap make` command:

```
# vxsnap -g mydg make source=myvol/snapvol=snap1myvol\
/syncing=off
```

- 2 Use `fsck` (or some utility appropriate for the application running on the volume) to clean the temporary volume’s contents. For example, you can use this command with a VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/snapshot
```

The specified device must have a valid entry in the `/etc/vfstab` file.

- 3 To backup the data in the snapshot, use an appropriate utility or operating system command to copy the contents of the snapshot to tape, or to some other backup medium.
- 4 You now have the following options:
 - Refresh the contents of the snapshot. This creates a new point-in-time image of the original volume ready for another backup. If synchronization was already in progress on the snapshot, this operation may result in large portions of the snapshot having to be resynchronized.

See [“Refreshing an instant space-optimized snapshot”](#) on page 616.

- Reattach some or all of the plexes of the snapshot volume with the original volume.
See [“Reattaching an instant full-sized or plex break-off snapshot”](#) on page 616.
- Restore the contents of the original volume from the snapshot volume. You can choose whether none, a subset, or all of the plexes of the snapshot volume are returned to the original volume as a result of the operation.
See [“Restoring a volume from an instant space-optimized snapshot”](#) on page 618.
- Dissociate the snapshot volume entirely from the original volume. This may be useful if you want to use the copy for other purposes such as testing or report generation. If desired, you can delete the dissociated volume.
See [“Dissociating an instant snapshot”](#) on page 618.
- If the snapshot is part of a snapshot hierarchy, you can also choose to split this hierarchy from its parent volumes.
See [“Splitting an instant snapshot hierarchy”](#) on page 619.

Creating and managing third-mirror break-off snapshots

Break-off snapshots are suitable for write-intensive volumes, such as database redo logs.

To turn one or more existing plexes in a volume into a break-off instant snapshot volume, the volume must be a non-layered volume with a `mirror` or `mirror-stripe` layout, or a RAID-5 volume that you have converted to a special layered volume and then mirrored. The plexes in a volume with a `stripe-mirror` layout are mirrored at the subvolume level, and cannot be broken off.

The attributes for a snapshot are specified as a tuple to the `vxsnap make` command. This command accepts multiple tuples. One tuple is required for each snapshot that is being created. Each element of a tuple is separated from the next by a slash character (/). Tuples are separated by white space.

To create and manage a third-mirror break-off snapshot

- 1 To create the snapshot, you can either take some of the existing `ACTIVE` plexes in the volume, or you can use the following command to add new snapshot mirrors to the volume:

```
# vxsnap [-b] [-g diskgroup] addmir volume [nmirror=N] \  
[alloc=storage_attributes]
```

By default, the `vxsnap addmir` command adds one snapshot mirror to a volume unless you use the `nmirror` attribute to specify a different number of mirrors. The mirrors remain in the `SNAPATT` state until they are fully synchronized. The `-b` option can be used to perform the synchronization in the background. Once synchronized, the mirrors are placed in the `SNAPDONE` state.

For example, the following command adds 2 mirrors to the volume, `vol1`, on disks `mydg10` and `mydg11`:

```
# vxsnap -g mydg addmir vol1 nmirror=2 alloc=mydg10,mydg11
```

If you specify the `-b` option to the `vxsnap addmir` command, you can use the `vxsnap snapwait` command to wait for synchronization of the snapshot plexes to complete, as shown in this example:

```
# vxsnap -g mydg snapwait vol1 nmirror=2
```

- 2 To create a third-mirror break-off snapshot, use the following form of the `vxsnap` `make` command.

```
# vxsnap [-g diskgroup] make source=volume[/newvol=snapvol]\
{/plex=plex1[,plex2,...] | /nmirror=number} }
```

Either of the following attributes may be specified to create the new snapshot volume, *snapvol*, by breaking off one or more existing plexes in the original volume:

<code>plex</code>	Specifies the plexes in the existing volume that are to be broken off.
<code>nmirror</code>	Specifies how many plexes are to be broken off. This attribute can only be used with plexes that are in the <code>SNAPDONE</code> state. (Such plexes could have been added to the volume by using the <code>vxsnap addmir</code> command.)

Snapshots that are created from one or more `ACTIVE` or `SNAPDONE` plexes in the volume are already synchronized by definition.

For backup purposes, a snapshot volume with one plex should be sufficient.

For example, to create the instant snapshot volume, `snap2myvol`, of the volume, `myvol`, in the disk group, `mydg`, from a single existing plex in the volume, use the following command:

```
# vxsnap -g mydg make source=myvol/newvol=snap2myvol/nmirror=1
```

The next example shows how to create a mirrored snapshot from two existing plexes in the volume:

```
# vxsnap -g mydg make source=myvol/newvol=snap2myvol/plex=myvol-03,myvol-04
```

- 3 Use `fsck` (or some utility appropriate for the application running on the volume) to clean the temporary volume's contents. For example, you can use this command with a VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/snapshot
```

The specified device must have a valid entry in the `/etc/vfstab` file.

- 4 To backup the data in the snapshot, use an appropriate utility or operating system command to copy the contents of the snapshot to tape, or to some other backup medium.
- 5 You now have the following options:

- Refresh the contents of the snapshot. This creates a new point-in-time image of the original volume ready for another backup. If synchronization was already in progress on the snapshot, this operation may result in large portions of the snapshot having to be resynchronized.
See [“Refreshing an instant space-optimized snapshot”](#) on page 616.
- Reattach some or all of the plexes of the snapshot volume with the original volume.
See [“Reattaching an instant full-sized or plex break-off snapshot”](#) on page 616.
- Restore the contents of the original volume from the snapshot volume. You can choose whether none, a subset, or all of the plexes of the snapshot volume are returned to the original volume as a result of the operation.
See [“Restoring a volume from an instant space-optimized snapshot”](#) on page 618.
- Dissociate the snapshot volume entirely from the original volume. This may be useful if you want to use the copy for other purposes such as testing or report generation. If desired, you can delete the dissociated volume.
See [“Dissociating an instant snapshot”](#) on page 618.
- If the snapshot is part of a snapshot hierarchy, you can also choose to split this hierarchy from its parent volumes.
See [“Splitting an instant snapshot hierarchy”](#) on page 619.

Creating and managing linked break-off snapshot volumes

Linked break-off snapshots are suitable for write-intensive volumes. Specifically, they are used for off-host processing, because the snapshot could be in a different disk group to start with and could avoid disk group split/join operations

For linked break-off snapshots, you must prepare a volume that is to be used as the snapshot volume. This must be the same size as the volume for which the snapshot is being created, and it must also have the same region size.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

The attributes for a snapshot are specified as a tuple to the `vxsnap make` command. This command accepts multiple tuples. One tuple is required for each snapshot that is being created. Each element of a tuple is separated from the next by a slash character (/). Tuples are separated by white space.

To create and manage a linked break-off snapshot

- 1 Use the following command to link the prepared snapshot volume, *snapvol*, to the data volume:

```
# vxsnap [-g diskgroup] [-b] addmir volume mirvol=snapvol \
    [mirdg=snapdg]
```

The optional *mirdg* attribute can be used to specify the snapshot volume's current disk group, *snapdg*. The *-b* option can be used to perform the synchronization in the background. If the *-b* option is not specified, the command does not return until the link becomes *ACTIVE*.

For example, the following command links the prepared volume, *prepsnap*, in the disk group, *mysnapdg*, to the volume, *vol1*, in the disk group, *mydg*:

```
# vxsnap -g mydg -b addmir vol1 mirvol=prepsnap mirdg=mysnapdg
```

If the *-b* option is specified, you can use the *vxsnap snapwait* command to wait for the synchronization of the linked snapshot volume to complete, as shown in this example:

```
# vxsnap -g mydg snapwait vol1 mirvol=prepsnap mirdg=mysnapvoldg
```

- 2 To create a linked break-off snapshot, use the following form of the *vxsnap make* command.

```
# vxsnap [-g diskgroup] make    source=volume/snapvol=snapvol\
    [/snapdg=snapdiskgroup]
```

The *snapdg* attribute must be used to specify the snapshot volume's disk group if this is different from that of the data volume.

For example, to use the prepared volume, *prepsnap*, as the snapshot for the volume, *vol1*, in the disk group, *mydg*, use the following command:

```
# vxsnap -g mydg make source=vol1/snapvol=prepsnap/snapdg=mysnapdg
```

- 3 Use *fsck* (or some utility appropriate for the application running on the volume) to clean the temporary volume's contents. For example, you can use this command with a VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/snapshot
```

The specified device must have a valid entry in the */etc/vfstab* file.

- 4 To backup the data in the snapshot, use an appropriate utility or operating system command to copy the contents of the snapshot to tape, or to some other backup medium.
- 5 You now have the following options:
 - Refresh the contents of the snapshot. This creates a new point-in-time image of the original volume ready for another backup. If synchronization was already in progress on the snapshot, this operation may result in large portions of the snapshot having to be resynchronized.
See [“Refreshing an instant space-optimized snapshot”](#) on page 616.
 - Reattach the snapshot volume with the original volume.
See [“Reattaching a linked break-off snapshot volume”](#) on page 617.
 - Dissociate the snapshot volume entirely from the original volume. This may be useful if you want to use the copy for other purposes such as testing or report generation. If desired, you can delete the dissociated volume.
See [“Dissociating an instant snapshot”](#) on page 618.
 - If the snapshot is part of a snapshot hierarchy, you can also choose to split this hierarchy from its parent volumes.
See [“Splitting an instant snapshot hierarchy”](#) on page 619.

Creating multiple instant snapshots

You can create multiple instant snapshots for all volumes that form a consistent group. The `vxsnap make` command accepts multiple tuples that define the source and snapshot volumes names as their arguments. For example, to create three instant snapshots, each with the same redundancy, from specified storage, the following form of the command can be used:

```
# vxsnap [-g diskgroup] make source=vol1/snapvol=snapvol1\  
source=vol2/snapvol=snapvol2 source=vol3/snapvol=snapvol3
```

The snapshot volumes (*snapvol1*, *snapvol2* and so on) must have been prepared in advance.

See [“Creating a volume for use as a full-sized instant or linked break-off snapshot”](#) on page 600.

The specified source volumes (*vol1*, *vol2* and so on) may be the same volume or they can be different volumes.

If all the snapshots are to be space-optimized and to share the same cache, the following form of the command can be used:

```
# vxsnap [-g diskgroup] make \
  source=vol1/newvol=snapvol1/cache=cacheobj \
  source=vol2/newvol=snapvol2/cache=cacheobj \
  source=vol3/newvol=snapvol3/cache=cacheobj \
  [alloc=storage_attributes]
```

The `vxsnap make` command also allows the snapshots to be of different types, have different redundancy, and be configured from different storage, as shown here:

```
# vxsnap [-g diskgroup] make source=vol1/snapvol=snapvol1 \
  source=vol2[/newvol=snapvol2]/cache=cacheobj \
  [/alloc=storage_attributes2] [/nmirror=number2]
  source=vol3/newvol=snapvol3 [/alloc=storage_attributes3] \
  /nmirror=number3
```

In this example, *snapvol1* is a full-sized snapshot that uses a prepared volume, *snapvol2* is a space-optimized snapshot that uses a prepared cache, and *snapvol3* is a break-off full-sized snapshot that is formed from plexes of the original volume.

An example of where you might want to create mixed types of snapshots at the same time is when taking snapshots of volumes containing database redo logs and database tables:

```
# vxsnap -g mydg make \
  source=logv1/newvol=snplogv1/drl=sequential/nmirror=1 \
  source=logv2/newvol=snplogv2/drl=sequential/nmirror=1 \
  source=datav1/newvol=snpdatav1/cache=mydgcobj/drl=on \
  source=datav2/newvol=snpdatav2/cache=mydgcobj/drl=on
```

In this example, sequential DRL is enabled for the snapshots of the redo log volumes, and normal DRL is applied to the snapshots of the volumes that contain the database tables. The two space-optimized snapshots are configured to share the same cache object in the disk group. Also note that break-off snapshots are used for the redo logs as such volumes are write intensive.

Creating instant snapshots of volume sets

Volume set names can be used in place of volume names with the following `vxsnap` operations on instant snapshots: `addmir`, `dis`, `make`, `prepare`, `reattach`, `refresh`, `restore`, `rmmir`, `split`, `syncpause`, `syncresume`, `syncstart`, `syncstop`, `syncwait`, and `unprepare`.

The procedure for creating an instant snapshot of a volume set is the same as that for a standalone volume. However, there are certain restrictions if a full-sized instant snapshot is to be created from a prepared volume set. A full-sized instant snapshot

of a volume set must itself be a volume set with the same number of volumes, and the same volume sizes and index numbers as the parent. For example, if a volume set contains three volumes with sizes 1GB, 2GB and 3GB, and indexes 0, 1 and 2 respectively, then the snapshot volume set must have three volumes with the same sizes matched to the same set of index numbers. The corresponding volumes in the parent and snapshot volume sets are also subject to the same restrictions as apply between standalone volumes and their snapshots.

You can use the `vxvset list` command to verify that the volume sets have identical characteristics as shown in this example:

```
# vxvset -g mydg list vset1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
vol_0	0	204800	ENABLED	-
vol_1	1	409600	ENABLED	-
vol_2	2	614400	ENABLED	-

```
# vxvset -g mydg list snapvset1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
svol_0	0	204800	ENABLED	-
svol_1	1	409600	ENABLED	-
svol_2	2	614400	ENABLED	-

A full-sized instant snapshot of a volume set can be created using a prepared volume set in which each volume is the same size as the corresponding volume in the parent volume set. Alternatively, you can use the `nmirrors` attribute to specify the number of plexes that are to be broken off provided that sufficient plexes exist for each volume in the volume set.

The following example shows how to prepare a source volume set, `vset1`, and an identical volume set, `snapvset1`, which is then used to create the snapshot:

```
# vxsnap -g mydg prepare vset1
# vxsnap -g mydg prepare snapvset1
# vxsnap -g mydg make source=vset1/snapvol=snapvset1
```

To create a full-sized third-mirror break-off snapshot, you must ensure that each volume in the source volume set contains sufficient plexes. The following example shows how to achieve this by using the `vxsnap` command to add the required number of plexes before breaking off the snapshot:

```
# vxsnap -g mydg prepare vset2
# vxsnap -g mydg addmir vset2 nmirror=1
# vxsnap -g mydg make source=vset2/newvol=snapvset2/nmirror=1
```

See [“Adding snapshot mirrors to a volume”](#) on page 614.

To create a space-optimized instant snapshot of a volume set, the commands are again identical to those for a standalone volume as shown in these examples:

```
# vxsnap -g mydg prepare vset3
# vxsnap -g mydg make source=vset3/newvol=snapvset3/cachesize=20m

# vxsnap -g mydg prepare vset4
# vxsnap -g mydg make source=vset4/newvol=snapvset4/cache=mycobj
```

Here a new cache object is created for the volume set, `vset3`, and an existing cache object, `mycobj`, is used for `vset4`.

Adding snapshot mirrors to a volume

If you are going to create a full-sized break-off snapshot volume, you can use the following command to add new snapshot mirrors to a volume:

```
# vxsnap [-b] [-g diskgroup] addmir volume|volume_set \
    [nmirror=N] [alloc=storage_attributes]
```

The volume must have been prepared using the `vxsnap prepare` command.

If a volume set name is specified instead of a volume, the specified number of plexes is added to each volume in the volume set.

By default, the `vxsnap addmir` command adds one snapshot mirror to a volume unless you use the `nmirror` attribute to specify a different number of mirrors. The mirrors remain in the `SNAPATT` state until they are fully synchronized. The `-b` option can be used to perform the synchronization in the background. Once synchronized, the mirrors are placed in the `SNAPDONE` state.

For example, the following command adds 2 mirrors to the volume, `vol1`, on disks `mydg10` and `mydg11`:

```
# vxsnap -g mydg addmir vol1 nmirror=2 alloc=mydg10,mydg11
```

This command is similar in usage to the `vxassist snapshot` command, and supports the traditional third-mirror break-off snapshot model. As such, it does not provide an instant snapshot capability.

Once you have added one or more snapshot mirrors to a volume, you can use the `vxsnap make` command with either the `nmirror` attribute or the `plex` attribute to create the snapshot volumes.

Removing a snapshot mirror

To remove a single snapshot mirror from a volume, use this command:

```
# vxsnap [-g diskgroup] rmmir volume|volume_set
```

For example, the following command removes a snapshot mirror from the volume, `voll`:

```
# vxsnap -g mydg rmmir voll
```

This command is similar in usage to the `vxassist snapabort` command.

If a volume set name is specified instead of a volume, a mirror is removed from each volume in the volume set.

Removing a linked break-off snapshot volume

To remove a linked break-off snapshot volume from a volume, use this command:

```
# vxsnap [-g diskgroup] rmmir volume|volume_set mirvol=snapvol \  
[mirdg=snapdiskgroup]
```

The `mirvol` and optional `mirdg` attributes specify the snapshot volume, *snapvol*, and its disk group, *snapdiskgroup*. For example, the following command removes a linked snapshot volume, `prepsnap`, from the volume, `voll`:

```
# vxsnap -g mydg rmmir voll mirvol=prepsnap mirdg=mysnapdg
```

Adding a snapshot to a cascaded snapshot hierarchy

To create a snapshot and push it onto a snapshot hierarchy between the original volume and an existing snapshot volume, specify the name of the existing snapshot volume as the value of the `infrontof` attribute to the `vxsnap make` command. The following example shows how to place the space-optimized snapshot, `thurs_bu`, of the volume, `dbvol`, in front of the earlier snapshot, `wed_bu`:

```
# vxsnap -g dbdg make source=dbvol/newvol=thurs_bu/  
infrontof=wed_bu/cache=dbdgcache
```

Similarly, the next snapshot that is taken, `fri_bu`, is placed in front of `thurs_bu`:

```
# vxsnap -g dbdg make source=dbvol/newvol=fri_bu/\
infrontof=thurs_bu/cache=dbdgcache
```

See [“Controlling instant snapshot synchronization”](#) on page 622.

Refreshing an instant space-optimized snapshot

Refreshing an instant snapshot replaces it with another point-in-time copy of a parent volume. To refresh one or more snapshots and make them immediately available for use, use the following command:

```
# vxsnap [-g diskgroup] refresh snapvolume|snapvolume_set \
[source=volume|volume_set] [snapvol2 [source=vol2]...] \
```

If the source volume is not specified, the immediate parent of the snapshot is used.

Warning: The snapshot that is being refreshed must not be open to any application. For example, any file system configured on the volume must first be unmounted.

Reattaching an instant full-sized or plex break-off snapshot

Using the following command, some or all plexes of an instant snapshot may be reattached to the specified original volume, or to a source volume in the snapshot hierarchy above the snapshot volume:

```
# vxsnap [-g diskgroup] reattach snapvolume|snapvolume_set \
source=volume|volume_set [nmirror=number]
```

By default, all the plexes are reattached, which results in the removal of the snapshot. If required, the number of plexes to be reattached may be specified as the value assigned to the `nmirror` attribute.

Warning: The snapshot that is being reattached must not be open to any application. For example, any file system configured on the snapshot volume must first be unmounted.

It is possible to reattach a volume to an unrelated volume provided that their volume sizes and region sizes are compatible.

For example the following command reattaches one plex from the snapshot volume, `snapmyvol`, to the volume, `myvol`:

```
# vxsnap -g mydg reattach snapmyvol source=myvol nmirror=1
```

While the reattached plexes are being resynchronized from the data in the parent volume, they remain in the `SNAPTMP` state. After resynchronization is complete, the plexes are placed in the `SNAPDONE` state. You can use the `vxsnap snapwait` command (but not `vxsnap syncwait`) to wait for the resynchronization of the reattached plexes to complete, as shown here:

```
# vxsnap -g mydg snapwait myvol nmirror=1
```

If the volume and its snapshot have both been resized (to an identical smaller or larger size) before performing the reattachment, a fast resynchronization can still be performed. A full resynchronization is not required. Instant snap DCO volumes are resized proportionately when the associated data volume is resized. For version 0 DCO volumes, the FastResync maps stay the same size, but the region size is recalculated, and the locations of the dirty bits in the existing maps are adjusted. In both cases, new regions are marked as dirty in the maps.

Reattaching a linked break-off snapshot volume

Unlike other types of snapshot, the reattachment operation for linked break-off snapshot volumes does not return the plexes of the snapshot volume to the parent volume. The link relationship is re-established that makes the snapshot volume a mirror of the parent volume, and this allows the snapshot data to be resynchronized.

To reattach a linked break-off snapshot volume, use the following form of the `vxsnap reattach` command:

```
# vxsnap [-g snapdiskgroup] reattach snapvolume|snapvolume_set \
    source=volume|volume_set [sourcedg=diskgroup]
```

The `sourcedg` attribute must be used to specify the data volume's disk group if this is different from the snapshot volume's disk group, `snapdiskgroup`.

Warning: The snapshot that is being reattached must not be open to any application. For example, any file system configured on the snapshot volume must first be unmounted.

It is possible to reattach a volume to an unrelated volume provided that their sizes and region sizes are compatible.

For example the following command reattaches the snapshot volume, `prepsnap`, in the disk group, `snapdg`, to the volume, `myvol`, in the disk group, `mydg`:

```
# vxsnap -g snapdg reattach prepsnap source=myvol sourcedg=mydg
```

After resynchronization of the snapshot volume is complete, the link is placed in the `ACTIVE` state. You can use the `vxsnap snapwait` command (but not `vxsnap syncwait`) to wait for the resynchronization of the reattached volume to complete, as shown here:

```
# vxsnap -g snapdg snapwait myvol mirvol=prepsnap
```

Restoring a volume from an instant space-optimized snapshot

It may sometimes be desirable to reinstate the contents of a volume from a backup or modified replica in a snapshot volume. The following command may be used to restore one or more volumes from the specified snapshots:

```
# vxsnap [-g diskgroup] restore volume|volume_set \
    source=snapvolume|snapvolume_set \
    [[volume2|volume_set2 \
    source=snapvolume2|snapvolume_set2]...] \
    [syncing=yes|no]
```

For a space-optimized instant snapshot, the cached data is used to recreate the contents of the specified volume. The space-optimized instant snapshot remains unchanged by the `restore` operation.

Warning: For this operation to succeed, the volume that is being restored and the snapshot volume must not be open to any application. For example, any file systems that are configured on either volume must first be unmounted.

It is not possible to restore a volume from an unrelated volume.

The following example demonstrates how to restore the volume, `myvol`, from the space-optimized snapshot, `snap3myvol`.

```
# vxsnap -g mydg restore myvol source=snap3myvol
```

Dissociating an instant snapshot

The following command breaks the association between a full-sized instant snapshot volume, `snapvol`, and its parent volume, so that the snapshot may be used as an independent volume:

```
# vxsnap [-f] [-g diskgroup] dis snapvolume|snapvolume_set
```

This operation fails if the snapshot, `snapvol`, has unsynchronized snapshots. If this happens, the dependent snapshots must be fully synchronized from `snapvol`. When

no dependent snapshots remain, *snapvol* may be dissociated. The snapshot hierarchy is then adopted by the parent volume of *snapvol*.

See [“Controlling instant snapshot synchronization”](#) on page 622.

See [“Removing an instant snapshot”](#) on page 619.

The following command dissociates the snapshot, *snap2myvol*, from its parent volume:

```
# vxsnap -g mydg dis snap2myvol
```

Warning: When applied to a volume set or to a component volume of a volume set, this operation can result in inconsistencies in the snapshot hierarchy in the case of a system crash or hardware failure. If the operation is applied to a volume set, the *-f* (force) option must be specified.

Removing an instant snapshot

When you have dissociated a full-sized instant snapshot, you can use the *vxedit* command to delete it altogether, as shown in this example:

```
# vxedit -g mydg -r rm snap2myvol
```

You can also use this command to remove a space-optimized instant snapshot from its cache.

See [“Removing a cache”](#) on page 626.

Splitting an instant snapshot hierarchy

Note: This operation is not supported for space-optimized instant snapshots.

The following command breaks the association between a snapshot hierarchy that has the snapshot volume, *snapvol*, at its head, and its parent volume, so that the snapshot hierarchy may be used independently of the parent volume:

```
# vxsnap [-f] [-g diskgroup] split snapvolume|snapvolume_set
```

The topmost snapshot volume in the hierarchy must have been fully synchronized for this command to succeed. Snapshots that are lower down in the hierarchy need not have been fully resynchronized.

See [“Controlling instant snapshot synchronization”](#) on page 622.

The following command splits the snapshot hierarchy under `snap2myvol` from its parent volume:

```
# vxsnap -g mydg split snap2myvol
```

Warning: When applied to a volume set or to a component volume of a volume set, this operation can result in inconsistencies in the snapshot hierarchy in the case of a system crash or hardware failure. If the operation is applied to a volume set, the `-f` (force) option must be specified.

Displaying instant snapshot information

The `vxsnap print` command may be used to display information about the snapshots that are associated with a volume.

```
# vxsnap [-g diskgroup] print [vol]
```

This command shows the percentage progress of the synchronization of a snapshot or volume. If no volume is specified, information about the snapshots for all the volumes in a disk group is displayed. The following example shows a volume, `vol1`, which has a full-sized snapshot, `snapvol1` whose contents have not been synchronized with `vol1`:

```
# vxsnap -g mydg print
```

NAME	SNAPOBJECT	TYPE	PARENT	SNAPSHOT	%DIRTY	%VALID
vol1	--	volume	--	--	--	100
	snapvol1_snp1	volume	--	snapvol1	1.30	--
snapvol1	vol1_snp1	volume	vol1	--	1.30	1.30

The `%DIRTY` value for `snapvol1` shows that its contents have changed by 1.30% when compared with the contents of `vol1`. As `snapvol1` has not been synchronized with `vol1`, the `%VALID` value is the same as the `%DIRTY` value. If the snapshot were partly synchronized, the `%VALID` value would lie between the `%DIRTY` value and 100%. If the snapshot were fully synchronized, the `%VALID` value would be 100%. The snapshot could then be made independent or moved into another disk group.

Additional information about the snapshots of volumes and volume sets can be obtained by using the `-n` option with the `vxsnap print` command:

```
# vxsnap [-g diskgroup] -n [-l] [-v] [-x] print [vol]
```

Alternatively, you can use the `vxsnap list` command, which is an alias for the `vxsnap -n print` command:


```
# vxsnap [-g diskgroup] [-l] [-v] [-x] list [vol]
```

The following output is an example of using this command on the disk group `dg1`:

```
# vxsnap -g dg -vx list
```

NAME	DG	OBJTYPE	SNAPTYPE	PARENT	PARENTDG	SNAPDATE	CHANGE_DATA	SYNCED_DATA
vol	dg1	vol	-	-	-	-	-	10G (100%)
svol1	dg2	vol	fullinst	vol	dg1	2006/2/1 12:29	20M (0.2%)	60M (0.6%)
svol2	dg1	vol	mirbrk	vol	dg1	2006/2/1 12:29	120M (1.2%)	10G (100%)
svol3	dg2	vol	volbrk	vol	dg1	2006/2/1 12:29	105M (1.1%)	10G (100%)
svol21	dg1	vol	spaceopt	svol2	dg1	2006/2/1 12:29	52M (0.5%)	52M (0.5%)
vol-02	dg1	plex	snapmir	vol	dg1	-	-	56M (0.6%)
mvol	dg2	vol	mirvol	vol	dg1	-	-	58M (0.6%)
vset1	dg1	vset	-	-	-	-	-	2G (100%)
v1	dg1	compvol	-	-	-	-	-	1G (100%)
v2	dg1	compvol	-	-	-	-	-	1G (100%)
svset1	dg1	vset	mirbrk	vset	dg1	2006/2/1 12:29	1G (50%)	2G (100%)
sv1	dg1	compvol	mirbrk	v1	dg1	2006/2/1 12:29	512M (50%)	1G (100%)
sv2	dg1	compvol	mirbrk	v2	dg1	2006/2/1 12:29	512M (50%)	1G (100%)
vol-03	dg1	plex	detmir	vol	dg1	-	20M (0.2%)	-
mvol2	dg2	vol	detvol	vol	dg1	-	20M (0.2%)	-

This shows that the volume `vol` has three full-sized snapshots, `svol1`, `svol2` and `svol3`, which are of types full-sized instant (`fullinst`), mirror break-off (`mirbrk`) and linked break-off (`volbrk`). It also has one snapshot `plex` (`snapmir`), `vol-02`, and one linked mirror volume (`mirvol`), `mvol`. The snapshot `svol2` itself has a space-optimized instant snapshot (`spaceopt`), `svol21`. There is also a volume set, `vset1`, with component volumes `v1` and `v2`. This volume set has a mirror break-off snapshot, `svset1`, with component volumes `sv1` and `sv2`. The last two entries show a detached plex, `vol-03`, and a detached mirror volume, `mvol2`, which have `vol` as their parent volume. These snapshot objects may have become detached due to an I/O error, or, in the case of the plex, by running the `vxplex det` command.

The `CHANGE_DATA` column shows the approximate difference between the current contents of the snapshot and its parent volume. This corresponds to the amount of data that would have to be resynchronized to make the contents the same again.

The `SYNCED_DATA` column shows the approximate progress of synchronization since the snapshot was taken.

The `-l` option can be used to obtain a longer form of the output listing instead of the tabular form.

The `-x` option expands the output to include the component volumes of volume sets.

See the `vxsnap(1M)` manual page for more information about using the `vxsnap print` and `vxsnap list` commands.

Controlling instant snapshot synchronization

Synchronization of the contents of a snapshot with its original volume is not possible for space-optimized instant snapshots.

By default, synchronization is enabled for the `vxsnap reattach`, `refresh` and `restore` operations on instant snapshots. Otherwise, synchronization is disabled unless you specify the `syncing=yes` attribute to the `vxsnap` command.

[Table 23-1](#) shows the commands that are provided for controlling the synchronization manually.

Table 23-1 Commands for controlling instant snapshot synchronization

Command	Description
<code>vxsnap [-g diskgroup] syncpause \</code> <code>vol vol_set</code>	Pause synchronization of a volume.
<code>vxsnap [-g diskgroup] syncresume \</code> <code>vol vol_set</code>	Resume synchronization of a volume.
<code>vxsnap [-b] [-g diskgroup] syncstart \</code> <code>vol vol_set</code>	Start synchronization of a volume. The <code>-b</code> option puts the operation in the background.
<code>vxsnap [-g diskgroup] syncstop \</code> <code>vol vol_set</code>	Stop synchronization of a volume.
<code>vxsnap [-g diskgroup] syncwait \</code> <code>vol vol_set</code>	Exit when synchronization of a volume is complete. An error is returned if the <code>vol</code> or <code>vol_set</code> is invalid (for example, it is a space-optimized snapshot), or if the <code>vol</code> or <code>vol_set</code> is not being synchronized. Note: You cannot use this command to wait for synchronization of reattached plexes to complete.

The commands that are shown in [Table 23-1](#) cannot be used to control the synchronization of linked break-off snapshots.

The `vxsnap snapwait` command is provided to wait for the link between new linked break-off snapshots to become ACTIVE, or for reattached snapshot plexes to reach the SNAPDONE state following resynchronization.

See [“Creating and managing linked break-off snapshot volumes”](#) on page 609.

See [“Reattaching an instant full-sized or plex break-off snapshot”](#) on page 616.

See [“Reattaching a linked break-off snapshot volume”](#) on page 617.

Improving the performance of snapshot synchronization

The following optional arguments to the `-o` option are provided to help optimize the performance of synchronization when using the `make`, `refresh`, `restore` and `syncstart` operations with full-sized instant snapshots:

<code>iosize=size</code>	Specifies the size of each I/O request that is used when synchronizing the regions of a volume. Specifying a larger size causes synchronization to complete sooner, but with greater impact on the performance of other processes that are accessing the volume. The default size of 1m (1MB) is suggested as the minimum value for high-performance array and controller hardware. The specified value is rounded to a multiple of the volume's region size.
<code>slow=iodelay</code>	Specifies the delay in milliseconds between synchronizing successive sets of regions as specified by the value of <code>iosize</code> . This can be used to change the impact of synchronization on system performance. The default value of <code>iodelay</code> is 0 milliseconds (no delay). Increasing this value slows down synchronization, and reduces the competition for I/O bandwidth with other processes that may be accessing the volume.

Options may be combined as shown in the following examples:

```
# vxsnap -g mydg -o iosize=2m,slow=100 make \  
    source=myvol/snapvol=snap2myvol/syncing=on  
  
# vxsnap -g mydg -o iosize=10m,slow=250 syncstart snap2myvol
```

Note: The `iosize` and `slow` parameters are not supported for space-optimized snapshots.

Listing the snapshots created on a cache

To list the space-optimized instant snapshots that have been created on a cache object, use the following command:

```
# vxcache [-g diskgroup] listvol cache_object
```

The snapshot names are printed as a space-separated list ordered by timestamp. If two or more snapshots have the same timestamp, these snapshots are sorted in order of decreasing size.

Tuning the autogrow attributes of a cache

The `highwatermark`, `autogrowby` and `maxautogrow` attributes determine how the VxVM cache daemon (`vxcached`) maintains the cache if the `autogrow` feature has been enabled and `vxcached` is running:

- When cache usage reaches the high watermark value, `highwatermark` (default value is 90 percent), `vxcached` grows the size of the cache volume by the value of `autogrowby` (default value is 20% of the size of the cache volume in blocks). The new required cache size cannot exceed the value of `maxautogrow` (default value is twice the size of the cache volume in blocks).
- When cache usage reaches the high watermark value, and the new required cache size would exceed the value of `maxautogrow`, `vxcached` deletes the oldest snapshot in the cache. If there are several snapshots with the same age, the largest of these is deleted.

If the `autogrow` feature has been disabled:

- When cache usage reaches the high watermark value, `vxcached` deletes the oldest snapshot in the cache. If there are several snapshots with the same age, the largest of these is deleted. If there is only a single snapshot, this snapshot is detached and marked as invalid.

Note: The `vxcached` daemon does not remove snapshots that are currently open, and it does not remove the last or only snapshot in the cache.

If the cache space becomes exhausted, the snapshot is detached and marked as invalid. If this happens, the snapshot is unrecoverable and must be removed. Enabling the `autogrow` feature on the cache helps to avoid this situation occurring. However, for very small caches (of the order of a few megabytes), it is possible for the cache to become exhausted before the system has time to respond and grow the cache. In such cases, you can increase the size of the cache manually.

Alternatively, you can use the `vxcache set` command to reduce the value of `highwatermark` as shown in this example:

```
# vxcache -g mydg set highwatermark=60 cobjmydg
```

You can use the `maxautogrow` attribute to limit the maximum size to which a cache can grow. To estimate this size, consider how much the contents of each source volume are likely to change between snapshot refreshes, and allow some additional space for contingency.

If necessary, you can use the `vxcache set` command to change other `autogrow` attribute values for a cache.

See the `vxcache(1M)` manual page.

Monitoring and displaying cache usage

You can use the `vxcache stat` command to display cache usage. For example, to see how much space is used and how much remains available in all cache objects in the diskgroup `mydg`, enter the following:

```
# vxcache -g mydg stat
```

Growing and shrinking a cache

You can use the `vxcache` command to increase the size of the cache volume that is associated with a cache object:

```
# vxcache [-g diskgroup] growcacheto cache_object  
size
```

For example, to increase the size of the cache volume associated with the cache object, `mycache`, to 2GB, you would use the following command:

```
# vxcache -g mydg growcacheto mycache 2g
```

To grow a cache by a specified amount, use the following form of the command shown here:

```
# vxcache [-g diskgroup] growcacheby cache_object  
size
```

For example, the following command increases the size of `mycache` by 1GB:

```
# vxcache -g mydg growcacheby mycache 1g
```

You can similarly use the `shrinkcacheby` and `shrinkcacheto` operations to reduce the size of a cache.

See the `vxcache(1M)` manual page.

Removing a cache

To remove a cache completely, including the cache object, its cache volume and all space-optimized snapshots that use the cache:

- 1 Run the following command to find out the names of the top-level snapshot volumes that are configured on the cache object:

```
# vxprint -g diskgroup -vne \
    "v_plex.pl_subdisk.sd_dm_name ~ /cache_object/"
```

where `cache_object` is the name of the cache object.

- 2 Remove all the top-level snapshots and their dependent snapshots (this can be done with a single command):

```
# vxedit -g diskgroup -r rm snapvol ...
```

where `snapvol` is the name of a top-level snapshot volume.

- 3 Stop the cache object:

```
# vxcache -g diskgroup stop cache_object
```

- 4 Finally, remove the cache object and its cache volume:

```
# vxedit -g diskgroup -r rm cache_object
```

Creating traditional third-mirror break-off snapshots

VxVM provides third-mirror break-off snapshot images of volume devices using `vxassist` and other commands.

To enhance the efficiency and usability of volume snapshots, turn on `FastResync`. If `Persistent FastResync` is required, you must associate a version 0 DCO with the volume.

See [“Adding a version 0 DCO and DCO volume”](#) on page 635.

A plex is required that is large enough to store the complete contents of the volume. Alternatively, you can use space-optimized instant snapshots.

The recommended approach to performing volume backup from the command line, or from a script, is to use the `vxsnap` command. The `vxassist snapstart`, `snapwait`, and `snapshot` commands are supported for backward compatibility.

The `vxassist snapshot` procedure consists of two steps:

- Run `vxassist snapstart` to create a snapshot mirror.
- Run `vxassist snapshot` to create a snapshot volume.

The `vxassist snapstart` step creates a write-only backup plex which gets attached to and synchronized with the volume. When synchronized with the volume, the backup plex is ready to be used as a `snapshot` mirror. The end of the update procedure is indicated by the new `snapshot` mirror changing its state to `SNAPDONE`. This change can be tracked by the `vxassist snapwait` task, which waits until at least one of the mirrors changes its state to `SNAPDONE`. If the attach process fails, the `snapshot` mirror is removed and its space is released.

Note: If the `snapstart` procedure is interrupted, the snapshot mirror is automatically removed when the volume is started.

Once the `snapshot` mirror is synchronized, it continues being updated until it is detached. You can then select a convenient time at which to create a `snapshot` volume as an image of the existing volume. You can also ask users to refrain from using the system during the brief time required to perform the `snapshot` (typically less than a minute). The amount of time involved in creating the `snapshot` mirror is long in contrast to the brief amount of time that it takes to create the `snapshot` volume.

The online backup procedure is completed by running the `vxassist snapshot` command on a volume with a `SNAPDONE` mirror. This task detaches the finished `snapshot` (which becomes a normal mirror), creates a new normal volume and attaches the `snapshot` mirror to the `snapshot` volume. The `snapshot` then becomes a normal, functioning volume and the state of the `snapshot` is set to `ACTIVE`.

To back up a volume using the vxassist command

- 1 Create a snapshot mirror for a volume using the following command:

```
# vxassist [-b] [-g diskgroup] snapstart [nmirror=N] volume
```

For example, to create a snapshot mirror of a volume called `voldef`, use the following command:

```
# vxassist [-g diskgroup] snapstart voldef
```

The `vxassist snapstart` task creates a write-only mirror, which is attached to and synchronized from the volume to be backed up.

By default, VxVM attempts to avoid placing snapshot mirrors on a disk that already holds any plexes of a data volume. However, this may be impossible if insufficient space is available in the disk group. In this case, VxVM uses any available space on other disks in the disk group. If the snapshot plexes are placed on disks which are used to hold the plexes of other volumes, this may cause problems when you subsequently attempt to move a snapshot volume into another disk group.

See [“Moving DCO volumes between disk groups”](#) on page 819.

To override the default storage allocation policy, you can use storage attributes to specify explicitly which disks to use for the snapshot plexes.

See [“Creating a volume on specific disks”](#) on page 204.

If you start `vxassist snapstart` in the background using the `-b` option, you can use the `vxassist snapwait` command to wait for the creation of the mirror to complete as shown here:

```
# vxassist [-g diskgroup] snapwait volume
```

If `vxassist snapstart` is not run in the background, it does not exit until the mirror has been synchronized with the volume. The mirror is then ready to be used as a plex of a snapshot volume. While attached to the original volume, its contents continue to be updated until you take the snapshot.

Use the `nmirror` attribute to create as many snapshot mirrors as you need for the snapshot volume. For a backup, you should usually only require the default of one.

It is also possible to make a snapshot plex from an existing plex in a volume.

See [“Converting a plex into a snapshot plex”](#) on page 630.

- 2 Choose a suitable time to create a snapshot. If possible, plan to take the snapshot at a time when users are accessing the volume as little as possible.

3 Create a snapshot volume using the following command:

```
# vxassist [-g diskgroup] snapshot [nmirror=N] volume snapshot
```

If required, use the `nmirror` attribute to specify the number of mirrors in the snapshot volume.

For example, to create a snapshot of `voldef`, use the following command:

```
# vxassist -g mydg snapshot voldef snapvoldef
```

The `vxassist snapshot` task detaches the finished snapshot mirror, creates a new volume, and attaches the snapshot mirror to it. This step should only take a few minutes. The snapshot volume, which reflects the original volume at the time of the snapshot, is now available for backing up, while the original volume continues to be available for applications and users.

If required, you can make snapshot volumes for several volumes in a disk group at the same time.

See [“Creating multiple snapshots with the vxassist command”](#) on page 631.

4 Use `fsck` (or some utility appropriate for the application running on the volume) to clean the temporary volume’s contents. For example, you can use this command with a VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/snapshot
```

The specified device must have a valid entry in the `/etc/vfstab` file.

5 If you require a backup of the data in the snapshot, use an appropriate utility or operating system command to copy the contents of the snapshot to tape, or to some other backup medium.

6 When the backup is complete, you have the following choices for what to do with the snapshot volume:

- Reattach some or all of the plexes of the snapshot volume with the original volume.
See [“Reattaching a snapshot volume”](#) on page 632.
- If FastResync was enabled on the volume before the snapshot was taken, this speeds resynchronization of the snapshot plexes before the backup cycle starts again at step 3.
- Dissociate the snapshot volume entirely from the original volume
See [“Dissociating a snapshot volume”](#) on page 633.
- This may be useful if you want to use the copy for other purposes such as testing or report generation.

- Remove the snapshot volume to save space with this command:

```
# vxedit [-g diskgroup] -rf rm snapshot
```

Dissociating or removing the snapshot volume loses the advantage of fast resynchronization if FastResync was enabled. If there are no further snapshot plexes available, any subsequent snapshots that you take require another complete copy of the original volume to be made.

Converting a plex into a snapshot plex

For a traditional, third-mirror break-off snapshot, you can convert an existing plex in a volume into a snapshot plex. Symantec recommends using the instant snapshot feature rather than converting a plex into a snapshot plex.

Note: A plex cannot be converted into a snapshot plex for layered volumes or for any volume that has an associated instant snap DCO volume.

In some circumstances, you may find it more convenient to convert an existing plex in a volume into a snapshot plex rather than running `vxassist snapstart`. For example, you may want to do this if you are short of disk space for creating the snapshot plex and the volume that you want to snapshot contains more than two plexes.

The procedure can also be used to speed up the creation of a snapshot volume when a mirrored volume is created with more than two plexes and `init=active` is specified.

It is advisable to retain at least two plexes in a volume to maintain data redundancy.

To convert an existing plex into a snapshot plex for a volume on which Persistent FastResync is enabled, use the following command:

```
# vxplex [-g diskgroup] -o dcoplex=dcologplex convert \  
state=SNAPDONE plex
```

`dcologplex` is the name of an existing DCO plex that is to be associated with the new snapshot plex. You can use the `vxprint` command to find out the name of the DCO volume.

See [“Adding a version 0 DCO and DCO volume”](#) on page 635.

For example, to make a snapshot plex from the plex `trivol-03` in the 3-plex volume `trivol`, you would use the following command:

```
# vxplex -o dcoplex=trivol_dco-03 convert state=SNAPDONE \
    trivial-03
```

Here the DCO plex `trivol_dco_03` is specified as the DCO plex for the new snapshot plex.

To convert an existing plex into a snapshot plex in the SNAPDONE state for a volume on which Non-Persistent FastResync is enabled, use the following command:

```
# vxplex [-g diskgroup] convert state=SNAPDONE plex
```

A converted plex is in the SNAPDONE state, and can be used immediately to create a snapshot volume.

Note: The last complete regular plex in a volume, an incomplete regular plex, or a dirty region logging (DRL) log plex cannot be converted into a snapshot plex.

See [“Third-mirror break-off snapshots”](#) on page 568.

Creating multiple snapshots with the vxassist command

To make it easier to create snapshots of several volumes at the same time, the snapshot option accepts more than one volume name as its argument, for example:

```
# vxassist [-g diskgroup] snapshot volume1
    volume2 ...
```

By default, the first snapshot volume is named *SNAP-volume*, and each subsequent snapshot is named *SNAPnumber-volume*, where *number* is a unique serial number, and *volume* is the name of the volume for which the snapshot is being taken. This default pattern can be overridden by using the option `-o name=pattern`, as described on the `vxassist(1M)` manual page. For example, the pattern `SNAP%v-%d` reverses the order of the *number* and *volume* components in the name.

To snapshot all the volumes in a single disk group, specify the option `-o allvols` to `vxassist`:

```
# vxassist -g diskgroup -o allvols snapshot
```

This operation requires that all `snapstart` operations are complete on the volumes. It fails if any of the volumes in the disk group do not have a complete snapshot plex in the SNAPDONE state.

Note: The `vxsnap` command provides similar functionality for creating multiple snapshots.

Reattaching a snapshot volume

The snapback operation merges a snapshot copy of a volume with the original volume. One or more snapshot plexes are detached from the snapshot volume and re-attached to the original volume. The snapshot volume is removed if all its snapshot plexes are snapped back. This task resynchronizes the data in the volume so that the plexes are consistent.

The snapback operation cannot be applied to RAID-5 volumes unless they have been converted to a special layered volume layout by the addition of a DCO and DCO volume.

See [“Adding a version 0 DCO and DCO volume”](#) on page 635.

To enhance the efficiency of the snapback operation, enable FastResync on the volume before taking the snapshot

To merge one snapshot plex with the original volume, use the following command:

```
# vxassist [-g diskgroup] snapback snapshot
```

where *snapshot* is the snapshot copy of the volume.

To merge all snapshot plexes in the snapshot volume with the original volume, use the following command:

```
# vxassist [-g diskgroup] -o allplexes snapback snapshot
```

To merge a specified number of plexes from the snapshot volume with the original volume, use the following command:

```
# vxassist [-g diskgroup] snapback nmirror=number  
      snapshot
```

Here the `nmirror` attribute specifies the number of mirrors in the snapshot volume that are to be re-attached.

Once the snapshot plexes have been reattached and their data resynchronized, they are ready to be used in another `snapshot` operation.

By default, the data in the original volume is used to update the snapshot plexes that have been re-attached. To copy the data from the replica volume instead, use the following command:

```
# vxassist [-g diskgroup] -o resyncfromreplica snapback snapshot
```

Warning: Always unmount the snapshot volume (if this is mounted) before performing a snapback. In addition, you must unmount the file system corresponding to the primary volume before using the `resyncfromreplica` option.

Adding plexes to a snapshot volume

If you want to retain the existing plexes in a snapshot volume after a snapback operation, you can create additional snapshot plexes that are to be used for the snapback.

To add plexes to a snapshot volume

- 1 Use the following `vxprint` commands to discover the names of the snapshot volume's data change object (DCO) and DCO volume:

```
# DCONAME=`vxprint [-g diskgroup] -F%dco_name snapshot`  
# DCOVOL=`vxprint [-g diskgroup] -F%log_vol $DCONAME`
```

- 2 Use the `vxassist mirror` command to create mirrors of the existing snapshot volume and its DCO volume:

```
# vxassist -g diskgroup mirror snapshot  
# vxassist -g diskgroup mirror $DCOVOL
```

The new plex in the DCO volume is required for use with the new data plex in the snapshot.

- 3 Use the `vxprint` command to find out the name of the additional snapshot plex:

```
# vxprint -g diskgroup snapshot
```

- 4 Use the `vxprint` command to find out the record ID of the additional DCO plex:

```
# vxprint -g diskgroup -F%rid $DCOVOL
```

- 5 Use the `vxedit` command to set the `dco_plex rid` field of the new data plex to the name of the new DCO plex:

```
# vxedit -g diskgroup set dco_plex rid=dco_plex rid new_plex
```

The new data plex is now ready to be used to perform a snapback operation.

Dissociating a snapshot volume

The link between a snapshot and its original volume can be permanently broken so that the snapshot volume becomes an independent volume. Use the following command to dissociate the snapshot volume, *snapshot*:

```
# vxassist snapclear snapshot
```

Displaying snapshot information

The `vxassist snapprint` command displays the associations between the original volumes and their respective replicas (snapshot copies):

```
# vxassist snapprint [volume]
```

Output from this command is shown in the following examples:

```
# vxassist -g mydg snapprint v1
```

V	NAME	USETYPE	LENGTH	
SS	SNAPOBJ	NAME	LENGTH	%DIRTY
DP	NAME	VOLUME	LENGTH	%DIRTY
v	v1	fsgen	20480	
ss	SNAP-v1_snp	SNAP-v1	20480	4
dp	v1-01	v1	20480	0
dp	v1-02	v1	20480	0
v	SNAP-v1	fsgen	20480	
ss	v1_snp	v1	20480	0

```
# vxassist -g mydg snapprint v2
```

V	NAME	USETYPE	LENGTH	
SS	SNAPOBJ	NAME	LENGTH	%DIRTY
DP	NAME	VOLUME	LENGTH	%DIRTY
v	v2	fsgen	20480	
ss	--	SNAP-v2	20480	0
dp	v2-01	v2	20480	0
v	SNAP-v2	fsgen	20480	
ss	--	v2	20480	0

In this example, Persistent FastResync is enabled on volume `v1`, and Non-Persistent FastResync on volume `v2`. Lines beginning with `v`, `dp` and `ss` indicate a volume, detached plex and snapshot plex respectively. The `%DIRTY` field indicates the percentage of a snapshot plex or detached plex that is dirty with respect to the original volume. Notice that no snap objects are associated with volume `v2` or with its snapshot volume `SNAP-v2`.

If a volume is specified, the `snapprint` command displays an error message if no FastResync maps are enabled for that volume.

Adding a version 0 DCO and DCO volume

To put Persistent FastResync into effect for a volume, a Data Change Object (DCO) and DCO volume must be associated with that volume. After you add a DCO object and DCO volume to a volume, you can enable Persistent FastResync on the volume.

Note: You need a FastResync license key to use the FastResync feature. Even if you do not have a license, you can configure a DCO object and DCO volume so that snap objects are associated with the original and snapshot volumes.

The procedure in this section describes adding a version 0 layout DCO. A version 0 DCO layout supports traditional (third-mirror break-off) snapshots that are administered with the `vxassist` command. A version 0 DCO layout does not support full-sized or space-optimized instant snapshots.

To add a DCO object and DCO volume to an existing volume

- 1 Ensure that the disk group containing the existing volume has at least disk group version 90. To check the version of a disk group:

```
# vxdg list diskgroup
```

If required, pgrade the disk group to the latest version:

```
# vxdg upgrade diskgroup
```

- 2 Turn off Non-Persistent FastResync on the original volume if it is currently enabled:

```
# vxvol [-g diskgroup] set fastresync=off volume
```

If you are uncertain about which volumes have Non-Persistent FastResync enabled, use the following command to obtain a listing of such volumes.

Note: The ! character is a special character in some shells. The following example shows how to escape it in a bash shell.

```
# vxprint [-g diskgroup] -F "%name" \  
-e "v_fastresync=on && \!v_hasdcolog"
```

- 3 Add a DCO and DCO volume to the existing volume (which may already have dirty region logging (DRL) enabled):

```
# vxassist [-g diskgroup] addlog volume logtype=dco \  
[ndcomirror=number] [dcolen=size] [storage_attributes]
```

For non-layered volumes, the default number of plexes in the mirrored DCO volume is equal to the lesser of the number of plexes in the data volume or 2. For layered volumes, the default number of DCO plexes is always 2. If required, use the `ndcomirror` attribute to specify a different number. It is recommended that you configure as many DCO plexes as there are existing data and snapshot plexes in the volume. For example, specify `ndcomirror=3` when adding a DCO to a 3-way mirrored volume.

The default size of each plex is 132 blocks. You can use the `dcolen` attribute to specify a different size. If specified, the size of the plex must be an integer multiple of 33 blocks from 33 up to a maximum of 2112 blocks.

You can specify `vxassist`-style storage attributes to define the disks that can or cannot be used for the plexes of the DCO volume.

See [“Specifying storage for version 0 DCO plexes”](#) on page 636.

Specifying storage for version 0 DCO plexes

If the disks that contain volumes and their snapshots are to be moved or split into different disk groups, the disks that contain their respective DCO plexes must be able to accompany them. By default, VxVM attempts to place version 0 DCO plexes on the same disks as the data plexes of the parent volume. However, this may be impossible if there is insufficient space available on those disks. In this case, VxVM uses any available space on other disks in the disk group. If the DCO plexes are

placed on disks which are used to hold the plexes of other volumes, this may cause problems when you subsequently attempt to move volumes into other disk groups.

You can use storage attributes to specify explicitly which disks to use for the DCO plexes. If possible, specify the same disks as those on which the volume is configured.

For example, to add a DCO object and DCO volume with plexes on `mydg05` and `mydg06`, and a plex size of 264 blocks to the volume, `myvol`, in the disk group, `mydg`, use the following command:

```
# vxassist -g mydg addlog myvol logtype=dcu dcolen=264 mydg05 mydg06
```

To view the details of the DCO object and DCO volume that are associated with a volume, use the `vxprint` command. The following is partial `vxprint` output for the volume named `vol1` (the `TUTIL0` and `PUTIL0` columns are omitted for clarity):

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	...
v	vol1	fsgen	ENABLED	1024	-	ACTIVE	
pl	vol1-01	vol1	ENABLED	1024	-	ACTIVE	
sd	disk01-01	vol1-01	ENABLED	1024	0	-	
pl	vol1-02	vol1	ENABLED	1024	-	ACTIVE	
sd	disk02-01	vol1-02	ENABLED	1024	0	-	
dc	vol1_dco	vol1	-	-	-	-	
v	vol1_dcl	gen	ENABLED	132	-	ACTIVE	
pl	vol1_dcl-01	vol1_dcl	ENABLED	132	-	ACTIVE	
sd	disk03-01	vol1_dcl-01	ENABLED	132	0	-	
pl	vol1_dcl-02	vol1_dcl	ENABLED	132	-	ACTIVE	
sd	disk04-01	vol1_dcl-02	ENABLED	132	0	-	

In this output, the DCO object is shown as `vol1_dco`, and the DCO volume as `vol1_dcl` with 2 plexes, `vol1_dcl-01` and `vol1_dcl-02`.

If required, you can use the `vxassist move` command to relocate DCO plexes to different disks. For example, the following command moves the plexes of the DCO volume, `vol1_dcl`, for volume `vol1` from `disk03` and `disk04` to `disk07` and `disk08`.

Note: The `!` character is a special character in some shells. The following example shows how to escape it in a bash shell.

```
# vxassist -g mydg move vol1_dcl \!disk03 \!disk04 disk07 disk08
```

See [“Moving DCO volumes between disk groups”](#) on page 819.

See the `vxassist(1M)` manual page.

Removing a version 0 DCO and DCO volume

To dissociate a version 0 DCO object, DCO volume and any snap objects from a volume, use the following command:

```
# vxassist [-g diskgroup] remove log volume logtype=dcv
```

This completely removes the DCO object, DCO volume and any snap objects. It also has the effect of disabling FastResync for the volume.

Alternatively, you can use the `vxdco` command to the same effect:

```
# vxdco [-g diskgroup] [-o rm] dis dco_obj
```

The default name of the DCO object, `dco_obj`, for a volume is usually formed by appending the string `_dco` to the name of the parent volume. To find out the name of the associated DCO object, use the `vxprint` command on the volume.

To dissociate, but not remove, the DCO object, DCO volume and any snap objects from the volume, `myvol`, in the disk group, `mydg`, use the following command:

```
# vxdco -g mydg dis myvol_dco
```

This form of the command dissociates the DCO object from the volume but does not destroy it or the DCO volume. If the `-o rm` option is specified, the DCO object, DCO volume and its plexes, and any snap objects are also removed.

Warning: Dissociating a DCO and DCO volume disables Persistent FastResync on the volume. A full resynchronization of any remaining snapshots is required when they are snapped back.

See the `vxassist(1M)` manual page.

See the `vxdco(1M)` manual pages.

Reattaching a version 0 DCO and DCO volume

If a version 0 DCO object and DCO volume are not removed by specifying the `-o rm` option to `vxdco`, they can be reattached to the parent volume using the following command:

```
# vxdco [-g diskgroup] att volume dco_obj
```

For example, to reattach the DCO object, `myvol_dco`, to the volume, `myvol`, use the following command:

```
# vxdco -g mydg att myvol myvol_dco
```

See the `vxdco(1M)` manual page.

Administering Storage Checkpoints

This chapter includes the following topics:

- [About Storage Checkpoints](#)
- [Storage Checkpoint administration](#)
- [Storage Checkpoint space management considerations](#)
- [Restoring from a Storage Checkpoint](#)
- [Storage Checkpoint quotas](#)

About Storage Checkpoints

Veritas File System (VxFS) provides a Storage Checkpoint feature that quickly creates a persistent image of a file system at an exact point in time. Storage Checkpoints significantly reduce I/O overhead by identifying and maintaining only the file system blocks that have changed since the last Storage Checkpoint or backup via a copy-on-write technique.

See [“Copy-on-write”](#) on page 574.

Storage Checkpoints provide:

- Persistence through reboots and crashes.
- The ability for data to be immediately writeable by preserving the file system metadata, the directory hierarchy, and user data.

Storage Checkpoints are actually data objects that are managed and controlled by the file system. You can create, remove, and rename Storage Checkpoints because they are data objects with associated names.

See [“How a Storage Checkpoint works”](#) on page 572.

Unlike a disk-based mirroring technology that requires a separate storage space, Storage Checkpoints minimize the use of disk space by using a Storage Checkpoint within the same free space available to the file system.

After you create a Storage Checkpoint of a mounted file system, you can also continue to create, remove, and update files on the file system without affecting the logical image of the Storage Checkpoint. A Storage Checkpoint preserves not only the name space (directory hierarchy) of the file system, but also the user data as it existed at the moment the file system image was captured.

You can use a Storage Checkpoint in many ways. For example, you can use them to:

- Create a stable image of the file system that can be backed up to tape.
- Provide a mounted, on-disk backup of the file system so that end users can restore their own files in the event of accidental deletion. This is especially useful in a home directory, engineering, or email environment.
- Create a copy of an application's binaries before installing a patch to allow for rollback in case of problems.
- Create an on-disk backup of the file system in that can be used in addition to a traditional tape-based backup to provide faster backup and restore capabilities.
- Test new software on a point-in-time image of the primary fileset without jeopardizing the live data in the current primary fileset by mounting the Storage Checkpoints as writable.

Storage Checkpoint administration

Storage Checkpoint administrative operations require the `fsckptadm` utility.

See the `fsckptadm(1M)` manual page.

You can use the `fsckptadm` utility to create and remove Storage Checkpoints, change attributes, and ascertain statistical data. Every Storage Checkpoint has an associated name, which allows you to manage Storage Checkpoints; this name is limited to 127 characters and cannot contain a colon (:).

See [“Creating a Storage Checkpoint”](#) on page 642.

See [“Removing a Storage Checkpoint”](#) on page 643.

Storage Checkpoints require some space for metadata on the volume or set of volumes specified by the file system allocation policy or Storage Checkpoint allocation policy. The `fsckptadm` utility displays an error if the volume or set of volumes does not have enough free space to contain the metadata. You can roughly

approximate the amount of space required by the metadata using a method that depends on the disk layout version of the file system.

For disk layout Version 7, multiply the number of inodes by 1 byte, and add 1 or 2 megabytes to get the approximate amount of space required. You can determine the number of inodes with the `fsckptadm` utility.

Use the `fsvoladm` command to determine if the volume set has enough free space.

See the `fsvoladm(1M)` manual page.

The following example lists the volume sets and displays the storage sizes in human-friendly units:

```
# fsvoladm -H list /mnt0
devid    size      used      avail     name
0         20 GB     10 GB     10 GB     vol1
1         30 TB     10 TB     20 TB     vol2
```

Creating a Storage Checkpoint

The following example shows the creation of a nodata Storage Checkpoint named `thu_7pm` on `/mnt0` and lists all Storage Checkpoints of the `/mnt0` file system:

```
# fsckptadm -n create thu_7pm /mnt0
# fsckptadm list /mnt0
/mnt0
thu_7pm:
  ctime    = Thu 3 Mar 2005 7:00:17 PM PST
  mtime    = Thu 3 Mar 2005 7:00:17 PM PST
  flags    = nodata, largefiles
```

The following example shows the creation of a removable Storage Checkpoint named `thu_8pm` on `/mnt0` and lists all Storage Checkpoints of the `/mnt0` file system:

```
# fsckptadm -r create thu_8pm /mnt0
# fsckptadm list /mnt0
/mnt0
thu_8pm:
  ctime    = Thu 3 Mar 2005 8:00:19 PM PST
  mtime    = Thu 3 Mar 2005 8:00:19 PM PST
  flags    = largefiles, removable
thu_7pm:
  ctime    = Thu 3 Mar 2005 7:00:17 PM PST
  mtime    = Thu 3 Mar 2005 7:00:17 PM PST
  flags    = nodata, largefiles
```

Removing a Storage Checkpoint

You can delete a Storage Checkpoint by specifying the `remove` keyword of the `fsckptadm` command. Specifically, you can use either the synchronous or asynchronous method of removing a Storage Checkpoint; the asynchronous method is the default method. The synchronous method entirely removes the Storage Checkpoint and returns all of the blocks to the file system before completing the `fsckptadm` operation. The asynchronous method simply marks the Storage Checkpoint for removal and causes `fsckptadm` to return immediately. At a later time, an independent kernel thread completes the removal operation and releases the space used by the Storage Checkpoint.

In this example, `/mnt0` is a mounted VxFS file system with a Version 9 disk layout. This example shows the asynchronous removal of the Storage Checkpoint named `thu_8pm` and synchronous removal of the Storage Checkpoint named `thu_7pm`. This example also lists all the Storage Checkpoints remaining on the `/mnt0` file system after the specified Storage Checkpoint is removed:

```
# fsckptadm remove thu_8pm /mnt0
# fsckptadm list /mnt0
/mnt0
thu_7pm:
  ctime   = Thu 3 Mar 2005 7:00:17 PM PST
  mtime   = Thu 3 Mar 2005 7:00:17 PM PST
  flags   = nodata, largefiles
# fsckptadm -s remove thu_7pm /mnt0
# fsckptadm list /mnt0
/mnt0
```

Accessing a Storage Checkpoint

You can mount Storage Checkpoints using the `mount` command with the `mount` option `-o ckpt=ckpt_name`.

See the `mount_vxfs(1M)` manual page.

Observe the following rules when mounting Storage Checkpoints:

- Storage Checkpoints are mounted as read-only Storage Checkpoints by default. If you must write to a Storage Checkpoint, mount it using the `-o rw` option.
- If a Storage Checkpoint is currently mounted as a read-only Storage Checkpoint, you can remount it as a writable Storage Checkpoint using the `-o remount` option.
- To mount a Storage Checkpoint of a file system, first mount the file system itself.

- To unmount a file system, first unmount all of its Storage Checkpoints.

Warning: If you create a Storage Checkpoint for backup purposes, do not mount it as a writable Storage Checkpoint. You will lose the point-in-time image if you accidentally write to the Storage Checkpoint.

If older Storage Checkpoints already exist, write activity to a writable Storage Checkpoint can generate copy operations and increased space usage in the older Storage Checkpoints.

A Storage Checkpoint is mounted on a special pseudo device. This pseudo device does not exist in the system name space; the device is internally created by the system and used while the Storage Checkpoint is mounted. The pseudo device is removed after you unmount the Storage Checkpoint. A pseudo device name is formed by appending the Storage Checkpoint name to the file system device name using the colon character (:) as the separator.

For example, if a Storage Checkpoint named `may_23` belongs to the file system residing on the special device `/dev/vx/dsk/fsvol/vol1`, the Storage Checkpoint pseudo device name is:

```
/dev/vx/dsk/fsvol/vol1:may_23
```

- To mount the Storage Checkpoint named `may_23` as a read-only Storage Checkpoint on directory `/fsvol_may_23`, type:

```
# mount -F vxfs -o ckpt=may_23 /dev/vx/dsk/fsvol/vol1:may_23 \
/fsvol_may_23
```

Note: The `vol1` file system must already be mounted before the Storage Checkpoint can be mounted.

- To remount the Storage Checkpoint named `may_23` as a writable Storage Checkpoint, type:

```
# mount -F vxfs -o ckpt=may_23,remount,rw \
/dev/vx/dsk/fsvol/vol1:may_23 /fsvol_may_23
```

- To mount this Storage Checkpoint automatically when the system starts up, put the following entries in the `/etc/vfstab` file:

#device	device to	mount point	FS	fsck	mount	mount
#to mount	fsck		type	pass	at boot	options
/dev/vx/dsk/fsvol/	/dev/vx/dsk/	/fsvol	vxfs	1	yes	-


```

vol1                fsvol/vol1
/dev/vx/dsk/fsvol/ -          /fsvol_may_23 vxfs 0      yes      ckpt=may_23
vol1:may_23

```

- To mount a Storage Checkpoint of a cluster file system, you must also use the `-o cluster` option:

```

# mount -F vxfs -o cluster,ckpt=may_23 \
/dev/vx/dsk/fsvol/vol1:may_23 /fsvol_may_23

```

You can only mount a Storage Checkpoint cluster-wide if the file system that the Storage Checkpoint belongs to is also mounted cluster-wide. Similarly, you can only mount a Storage Checkpoint locally if the file system that the Storage Checkpoint belongs to is mounted locally.

You can unmount Storage Checkpoints using the `umount` command.

See the `umount_vxfs(1M)` manual page.

Storage Checkpoints can be unmounted by the mount point or pseudo device name:

```

# umount /fsvol_may_23
# umount /dev/vx/dsk/fsvol/vol1:may_23

```

Note: You do not need to run the `fsck` utility on Storage Checkpoint pseudo devices because pseudo devices are part of the actual file system.

Converting a data Storage Checkpoint to a nodata Storage Checkpoint

A nodata Storage Checkpoint does not contain actual file data. Instead, this type of Storage Checkpoint contains a collection of markers indicating the location of all the changed blocks since the Storage Checkpoint was created.

See [“Types of Storage Checkpoints”](#) on page 578.

You can use either the synchronous or asynchronous method to convert a data Storage Checkpoint to a nodata Storage Checkpoint; the asynchronous method is the default method. In a synchronous conversion, `fsckptadm` waits for all files to undergo the conversion process to “nodata” status before completing the operation. In an asynchronous conversion, `fsckptadm` returns immediately and marks the Storage Checkpoint as a nodata Storage Checkpoint even though the Storage Checkpoint’s data blocks are not immediately returned to the pool of free blocks in the file system. The Storage Checkpoint deallocates all of its file data blocks in the background and eventually returns them to the pool of free blocks in the file system.

If all of the older Storage Checkpoints in a file system are nodata Storage Checkpoints, use the synchronous method to convert a data Storage Checkpoint to a nodata Storage Checkpoint. If an older data Storage Checkpoint exists in the file system, use the asynchronous method to mark the Storage Checkpoint you want to convert for a delayed conversion. In this case, the actual conversion will continue to be delayed until the Storage Checkpoint becomes the oldest Storage Checkpoint in the file system, or all of the older Storage Checkpoints have been converted to nodata Storage Checkpoints.

Note: You cannot convert a nodata Storage Checkpoint to a data Storage Checkpoint because a nodata Storage Checkpoint only keeps track of the location of block changes and does not save the content of file data blocks.

Showing the difference between a data and a nodata Storage Checkpoint

The following example shows the difference between data Storage Checkpoints and nodata Storage Checkpoints.

Note: A nodata Storage Checkpoint does not contain actual file data.

To show the difference between Storage Checkpoints

- 1 Create a file system and mount it on `/mnt0`, as in the following example:

```
# mkfs -F vxfs /dev/vx/rdisk/dg1/test0

version 9 layout
134217728 sectors, 67108864 blocks of size 1024, log \

size 65536 blocks, largefiles supported
# mount -F vxfs /dev/vx/rdisk/dg1/test0 /mnt0
```

- 2 Create a small file with a known content, as in the following example:

```
# echo "hello, world" > /mnt0/file
```

- 3 Create a Storage Checkpoint and mount it on `/mnt0@5_30pm`, as in the following example:

```
# fsckptadm create ckpt@5_30pm /mnt0
# mkdir /mnt0@5_30pm
# mount -F vxfs -o ckpt=ckpt@5_30pm \
/dev/vx/dsk/dg1/test0:ckpt@5_30pm /mnt0@5_30pm
```

- 4 Examine the content of the original file and the Storage Checkpoint file:

```
# cat /mnt0/file
hello, world
# cat /mnt0@5_30pm/file
hello, world
```

- 5 Change the content of the original file:

```
# echo "goodbye" > /mnt0/file
```

- 6 Examine the content of the original file and the Storage Checkpoint file. The original file contains the latest data while the Storage Checkpoint file still contains the data at the time of the Storage Checkpoint creation:

```
# cat /mnt0/file
goodbye
# cat /mnt0@5_30pm/file
hello, world
```

- 7 Unmount the Storage Checkpoint, convert the Storage Checkpoint to a nodata Storage Checkpoint, and mount the Storage Checkpoint again:

```
# umount /mnt0@5_30pm
# fsckptadm -s set nodata ckpt@5_30pm /mnt0
# mount -F vxfs -o ckpt=ckpt@5_30pm \
/dev/vx/dsk/dg1/test0:ckpt@5_30pm /mnt0@5_30pm
```

- 8 Examine the content of both files. The original file must contain the latest data:

```
# cat /mnt0/file
goodbye
```

You can traverse and read the directories of the nodata Storage Checkpoint; however, the files contain no data, only markers to indicate which block of the file has been changed since the Storage Checkpoint was created:

```
# ls -l /mnt0@5_30pm/file
-rw-r--r--  1 root      other 13 Jul 13 17:13 \
# cat /mnt0@5_30pm/file
cat: input error on /mnt0@5_30pm/file: I/O error

# ls -l /mnt0@5_30pm/file
-rw-r--r--  1 root      other 13 Jul 13 17:13 \
# cat /mnt0@5_30pm/file
cat: /mnt0@5_30pm/file: I/O error
```

Converting multiple Storage Checkpoints

You can convert Storage Checkpoints to nodata Storage Checkpoints when dealing with older Storage Checkpoints on the same file system.

To convert multiple Storage Checkpoints

- 1 Create a file system and mount it on /mnt0:

```
# mkfs -F vxfs /dev/vx/rdisk/dg1/test0
version 9 layout
13417728 sectors, 67108864 blocks of size 1024, log \
size 65536 blocks largefiles supported
# mount -F vxfs /dev/vx/dsk/dg1/test0 /mnt0
```

- 2 Create four data Storage Checkpoints on this file system, note the order of creation, and list them:

```
# fsckptadm create oldest /mnt0
# fsckptadm create older /mnt0
# fsckptadm create old /mnt0
# fsckptadm create latest /mnt0
# fsckptadm list /mnt0
/mnt0
latest:
    ctime                = Mon 26 Jul 11:56:55 2004
    mtime                = Mon 26 Jul 11:56:55 2004
    flags                = largefiles
old:
    ctime                = Mon 26 Jul 11:56:51 2004
    mtime                = Mon 26 Jul 11:56:51 2004
    flags                = largefiles
older:
    ctime                = Mon 26 Jul 11:56:46 2004
    mtime                = Mon 26 Jul 11:56:46 2004
    flags                = largefiles
oldest:
    ctime                = Mon 26 Jul 11:56:41 2004
    mtime                = Mon 26 Jul 11:56:41 2004
    flags                = largefiles
```

- 3 Try to convert synchronously the latest Storage Checkpoint to a nodata Storage Checkpoint. The attempt will fail because the Storage Checkpoints older than the latest Storage Checkpoint are data Storage Checkpoints, namely the Storage Checkpoints old, older, and oldest:

```
# fsckptadm -s set nodata latest /mnt0
UX:vxfs fsckptadm: ERROR: V-3-24632: Storage Checkpoint
set failed on latest. File exists (17)
```

- 4 You can instead convert the `latest` Storage Checkpoint to a `nodata` Storage Checkpoint in a delayed or asynchronous manner.

```
# fsckptadm set nodata latest /mnt0
```

- 5 List the Storage Checkpoints, as in the following example. You will see that the `latest` Storage Checkpoint is marked for conversion in the future.

```
# fsckptadm list /mnt0
/mnt0
latest:
  ctime           = Mon 26 Jul 11:56:55 2004
  mtime           = Mon 26 Jul 11:56:55
  flags           = nodata, largefiles, delayed
old:
  ctime           = Mon 26 Jul 11:56:51 2004
  mtime           = Mon 26 Jul 11:56:51 2004
  flags           = largefiles
older:
  ctime           = Mon 26 Jul 11:56:46 2004
  mtime           = Mon 26 Jul 11:56:46 2004
  flags           = largefiles
oldest:
  ctime           = Mon 26 Jul 11:56:41 2004
  mtime           = Mon 26 Jul 11:56:41 2004
  flags           = largefiles
```

Creating a delayed nodata Storage Checkpoint

You can create a Storage Checkpoint as a delayed nodata Storage Checkpoint. The creation process detects the presence of the older data Storage Checkpoints and creates the Storage Checkpoint as a delayed nodata Storage Checkpoint. The following example procedure removes an existing Storage Checkpoint named `latest` and recreates the Storage Checkpoint as a delayed nodata Storage Checkpoint.

To create a delayed nodata Storage Checkpoint

1 Remove the latest Storage Checkpoint.

```
# fsckptadm remove latest /mnt0
# fsckptadm list /mnt0
/mnt0
old:
  ctime          = Mon 26 Jul 11:56:51 2004
  mtime          = Mon 26 Jul 11:56:51 2004
  flags          = largefiles
older:
  ctime          = Mon 26 Jul 11:56:46 2004
  mtime          = Mon 26 Jul 11:56:46 2004
  flags          = largefiles
oldest:
  ctime          = Mon 26 Jul 11:56:41 2004
  mtime          = Mon 26 Jul 11:56:41 2004
  flags          = largefiles
```

2 Recreate the latest Storage Checkpoint as a nodata Storage Checkpoint.

```
# fsckptadm -n create latest /mnt0
# fsckptadm list /mnt0
/mnt0
latest:
  ctime          = Mon 26 Jul 12:06:42 2004
  mtime          = Mon 26 Jul 12:06:42 2004
  flags          = nodata, largefiles, delayed
old:
  ctime          = Mon 26 Jul 11:56:51 2004
  mtime          = Mon 26 Jul 11:56:51 2004
  flags          = largefiles
older:
  ctime          = Mon 26 Jul 11:56:46 2004
  mtime          = Mon 26 Jul 11:56:46 2004
  flags          = largefiles
oldest:
  ctime          = Mon 26 Jul 11:56:41 2004
  mtime          = Mon 26 Jul 11:56:41 2004
  flags          = largefiles
```

- 3 Convert the `oldest` Storage Checkpoint to a `nodata` Storage Checkpoint because no older Storage Checkpoints exist that contain data in the file system.

Note: This step can be done synchronously.

```
# fsckptadm -s set nodata oldest /mnt0
# fsckptadm list /mnt0
/mnt0
latest:
  ctime           = Mon 26 Jul 12:06:42 2004
  mtime           = Mon 26 Jul 12:06:42 2004
  flags           = nodata, largefiles, delayed
old:
  ctime           = Mon 26 Jul 11:56:51 2004
  mtime           = Mon 26 Jul 11:56:51 2004
  flags           = largefiles
older:
  ctime           = Mon 26 Jul 11:56:46 2004
  mtime           = Mon 26 Jul 11:56:46 2004
  flags           = largefiles
oldest:
  ctime           = Mon 26 Jul 11:56:41 2004
  mtime           = Mon 26 Jul 11:56:41 2004
  flags           = nodata, largefiles
```


4 Remove the `older` and `old` Storage Checkpoints.

```
# fsckptadm remove older /mnt0
# fsckptadm remove old /mnt0
# fsckptadm list /mnt0
/mnt0
latest:
  ctime           = Mon 26 Jul 12:06:42 2004
  mtime           = Mon 26 Jul 12:06:42 2004
  flags           = nodata, largefiles
oldest:
  ctime           = Mon 26 Jul 11:56:41 2004
  mtime           = Mon 26 Jul 11:56:41 2004
  flags           = nodata, largefiles
```

Note: After you remove the `older` and `old` Storage Checkpoints, the `latest` Storage Checkpoint is automatically converted to a `nodata` Storage Checkpoint because the only remaining `older` Storage Checkpoint (`oldest`) is already a `nodata` Storage Checkpoint:

Enabling and disabling Storage Checkpoint visibility

You enable Storage Checkpoint visibility through the `ckptautomnt` mount option, which can be set to one of three values: `off`, `ro`, or `rw`. Because enabling Storage Checkpoint visibility prevents manual mounting of clones, the default value is `off`. Setting the option to `ro` causes all clones to be automounted as read-only, while `rw` causes all clones to be automounted as read/write.

If you take a Storage Checkpoint of an existing Storage Checkpoint (instead of the primary file set), the directory for the source Storage Checkpoint in `.checkpoint` functions as the mount point. For example, to take a Storage Checkpoint of the Storage Checkpoint `clone1` in a file system mounted on `/mnt`, use the following command:

```
# fsckptadm create clone2 /mnt/.checkpoint/clone1
```

By default, Veritas Storage Foundation (SF) does not make inode numbers unique. However, you can specify the `uniqueino` mount option to enable the use of unique 64-bit inode numbers. You cannot change this option during a remount.

The following example enables Storage Checkpoint visibility by causing all clones to be automounted as read/write:

```
# mount -F vxfs -o ckptautomnt=rw /dev/vx/dsk/dg1/vol1 /mnt1
```

Storage Checkpoint space management considerations

Several operations, such as removing or overwriting a file, can fail when a file system containing Storage Checkpoints runs out of space. If the system cannot allocate sufficient space, the operation will fail.

Database applications usually preallocate storage for their files and may not expect a write operation to fail. During user operations such as `create` or `mkdir`, if the file system runs out of space, removable Storage Checkpoints are deleted. This ensures that applications can continue without interruptions due to lack of disk space.

Non-removable Storage Checkpoints are not automatically removed under such `ENOSPC` conditions. Symantec recommends that you create only removable Storage Checkpoints. However, during certain administrative operations, such as using the `fsadm` command and creating a Storage Checkpoint with the `fsckptadm` command, even if the file system runs out of space, removable Storage Checkpoints are not deleted.

When the kernel automatically removes the Storage Checkpoints, it applies the following policies:

- Remove as few Storage Checkpoints as possible to complete the operation.
- Never select a non-removable Storage Checkpoint.
- Select a nodata Storage Checkpoint only when data Storage Checkpoints no longer exist.
- Remove the oldest Storage Checkpoint first.
- Remove a Storage Checkpoint even if it is mounted. New operations on such a removed Storage Checkpoint fail with the appropriate error codes.
- If the oldest Storage Checkpoint is non-removable, then the oldest removable Storage Checkpoint is selected for removal. In such a case, data might be required to be pushed to a non-removable Storage Checkpoint, which might fail and result in the file system getting marked for a FULLFSCK. To prevent this occurrence, Symantec recommends that you only create removable Storage Checkpoints.

Restoring from a Storage Checkpoint

Mountable data Storage Checkpoints on a consistent and undamaged file system can be used by backup and restore applications to restore either individual files or

an entire file system. Restoration from Storage Checkpoints can also help recover incorrectly modified files, but typically cannot recover from hardware damage or other file system integrity problems.

Note: For hardware or other integrity problems, Storage Checkpoints must be supplemented by backups from other media.

Files can be restored by copying the entire file from a mounted Storage Checkpoint back to the primary fileset. To restore an entire file system, you can designate a mountable data Storage Checkpoint as the primary fileset using the `fsckpt_restore` command.

See the `fsckpt_restore(1M)` manual page.

When using the `fsckpt_restore` command to restore a file system from a Storage Checkpoint, all changes made to that file system after that Storage Checkpoint's creation date are permanently lost. The only Storage Checkpoints and data preserved are those that were created at the same time, or before, the selected Storage Checkpoint's creation. The file system cannot be mounted at the time that `fsckpt_restore` is invoked.

Note: Individual files can also be restored very efficiently by applications using the `fsckpt_fmap(3)` library function to restore only modified portions of a file's data.

You can restore from a Storage Checkpoint only to a file system that has disk layout Version 6 or later.

Examples of restoring a file from a Storage Checkpoint

The following example restores a file, `MyFile.txt`, which resides in your home directory, from the Storage Checkpoint `CKPT1` to the device `/dev/vx/dsk/dg1/vol-01`. The mount point for the device is `/home`.

To restore a file from a Storage Checkpoint

- 1 Create the Storage Checkpoint `CKPT1` of `/home`.

```
$ fckptadm create CKPT1 /home
```

- 2 Mount Storage Checkpoint `CKPT1` on the directory `/home/checkpoints/mar_4`.

```
$ mount -F vxfs -o ckpt=CKPT1 /dev/vx/dsk/dg1/vol- \
01:CKPT1 /home/checkpoints/mar_4
```

- 3 Delete the file `MyFile.txt` from your home directory.

```
$ cd /home/users/me
$ rm MyFile.txt
```


- 4 Go to the `/home/checkpoints/mar_4/users/me` directory, which contains the image of your home directory.

```
$ cd /home/checkpoints/mar_4/users/me
$ ls -l
-rw-r--r--  1 me  staff  14910   Mar 4   17:09  MyFile.txt
```

- 5 Copy the file `MyFile.txt` to your home directory.

```
$ cp MyFile.txt /home/users/me
$ cd /home/users/me
$ ls -l
-rw-r--r--  1 me  staff  14910   Mar 4   18:21  MyFile.txt
```

The following example restores a file system from the Storage Checkpoint `CKPT3`. The filesets listed before the restoration show an unnamed root fileset and six Storage Checkpoints.



To restore a file system from a Storage Checkpoint

1 Run the `fsckpt_restore` command:

```
# fsckpt_restore -l /dev/vx/dsk/dg1/vol2
/dev/vx/dsk/dg1/vol2:
UNNAMED:
    ctime          = Thu 08 May 2004 06:28:26 PM PST
    mtime          = Thu 08 May 2004 06:28:26 PM PST
    flags          = largefiles, file system root
CKPT6:
    ctime          = Thu 08 May 2004 06:28:35 PM PST
    mtime          = Thu 08 May 2004 06:28:35 PM PST
    flags          = largefiles
CKPT5:
    ctime          = Thu 08 May 2004 06:28:34 PM PST
    mtime          = Thu 08 May 2004 06:28:34 PM PST
    flags          = largefiles, nomount
CKPT4:
    ctime          = Thu 08 May 2004 06:28:33 PM PST
    mtime          = Thu 08 May 2004 06:28:33 PM PST
    flags          = largefiles
CKPT3:
    ctime          = Thu 08 May 2004 06:28:36 PM PST
    mtime          = Thu 08 May 2004 06:28:36 PM PST
    flags          = largefiles
CKPT2:
    ctime          = Thu 08 May 2004 06:28:30 PM PST
    mtime          = Thu 08 May 2004 06:28:30 PM PST
    flags          = largefiles
CKPT1:
    ctime          = Thu 08 May 2004 06:28:29 PM PST
    mtime          = Thu 08 May 2004 06:28:29 PM PST
    flags          = nodata, largefiles
```

2 In this example, select the Storage Checkpoint CKPT3 as the new root fileset:

```
Select Storage Checkpoint for restore operation
or <Control/D> (EOF) to exit
or <Return> to list Storage Checkpoints: CKPT3
CKPT3:
    ctime          = Thu 08 May 2004 06:28:31 PM PST
    mtime          = Thu 08 May 2004 06:28:36 PM PST
    flags          = largefiles
UX:vxfs fsckpt_restore: WARNING: V-3-24640: Any file system
changes or Storage Checkpoints made after
Thu 08 May 2004 06:28:31 PM PST will be lost.
```


3 Type **y** to restore the file system from CKPT3:

```

Restore the file system from Storage Checkpoint CKPT3 ?
(ynq) y
(Yes)
UX:vxfs fsckpt_restore: INFO: V-3-23760: File system
restored from CKPT3

```

If the filesets are listed at this point, it shows that the former UNNAMED root fileset and CKPT6, CKPT5, and CKPT4 were removed, and that CKPT3 is now the primary fileset. CKPT3 is now the fileset that will be mounted by default.



4 Run the `fsckpt_restore` command:

```

# fsckpt_restore -l /dev/vx/dsk/dg1/vol2
/dev/vx/dsk/dg1/vol2:
CKPT3:
    ctime      = Thu 08 May 2004 06:28:31 PM PST
    mtime      = Thu 08 May 2004 06:28:36 PM PST
    flags      = largefiles, file system root
CKPT2:
    ctime      = Thu 08 May 2004 06:28:30 PM PST
    mtime      = Thu 08 May 2004 06:28:30 PM PST
    flags      = largefiles
CKPT1:
    ctime      = Thu 08 May 2004 06:28:29 PM PST
    mtime      = Thu 08 May 2004 06:28:29 PM PST
    flags      = nodata, largefiles
Select Storage Checkpoint for restore operation
or <Control/D> (EOF) to exit
or <Return> to list Storage Checkpoints:

```

Storage Checkpoint quotas

VxFS provides options to the `fsckptadm` command interface to administer Storage Checkpoint quotas. Storage Checkpoint quotas set the following limits on the amount of space used by all Storage Checkpoints of a primary file set:

hard limit	An absolute limit that cannot be exceeded. If a hard limit is exceeded, all further allocations on any of the Storage Checkpoints fail, but existing Storage Checkpoints are preserved.
soft limit	Must be lower than the hard limit. If a soft limit is exceeded, no new Storage Checkpoints can be created. The number of blocks used must return below the soft limit before more Storage Checkpoints can be created. An alert and console message are generated.

In case of a hard limit violation, various solutions are possible, enacted by specifying or not specifying the `-f` option for the `fsckptadm` utility.

See the `fsckptadm(1M)` manual page.

Specifying or not specifying the `-f` option has the following effects:

- If the `-f` option is not specified, one or many removable Storage Checkpoints are deleted to make space for the operation to succeed. This is the default solution.
- If the `-f` option is specified, all further allocations on any of the Storage Checkpoints fail, but existing Storage Checkpoints are preserved.

Note: Sometimes if a file is removed while it is opened by another process, the removal process is deferred until the last close. Because the removal of a file may trigger pushing data to a “downstream” Storage Checkpoint (that is, the next older Storage Checkpoint), a fileset hard limit quota violation may occur. In this scenario, the hard limit is relaxed to prevent an inode from being marked bad. This is also true for some asynchronous inode operations.

Administering FileSnaps

This chapter includes the following topics:

- [FileSnap creation](#)
- [Using FileSnaps](#)
- [Using FileSnaps to create point-in-time copies of files](#)
- [Comparison of the logical size output of the `fsadm -S shared`, `du`, and `df` commands](#)

FileSnap creation

A single thread creating FileSnaps of the same file can create over ten thousand snapshots per minute. FileSnaps can be used for fast provisioning of new virtual machines by cloning a virtual machine golden image, where the golden image is stored as a file in a VxFS file system or Veritas Storage Foundation Cluster File System (SFCFS) file system, which is used as a data store for a virtual environment.

FileSnap creation over Network File System

You can create a FileSnap over Network File System (NFS) by creating a hard link from an existing file to a new file with the extension “:snap:vxfs:”. For example, the following command creates a new file named `file1`, but instead of making `file1` a hard link of `file2`, `file1` will be a FileSnap so that the link count of `file2` will not change:

```
# ln file1 file2::snap:vxfs:
```

This is the equivalent of using the following command:

```
# vxfilesnap -p file1 file2
```

The new file has the same attributes as the old file and shares all of the old file's extents.

An application that uses this namespace extension should check if the file created has the namespace extension, such as `file1::snap:vxfs:` instead of `file1`. This indicates the namespace extension is not supported, either because the file system exported over NFS is not VxFS, the file system is an older version of VxFS, or the file system does not have a license for FileSnaps.

As with the `vxfilesnap` command, FileSnaps must be made within a single file set.

Using FileSnaps

[Table 25-1](#) provides a list of Veritas File System (VxFS) commands that enable you to administer FileSnaps.

Table 25-1

Command	Functionality
<code>fiostat</code>	The <code>fiostat</code> command has the <code>-S shared</code> option to display statistics for each interval. Otherwise, the command displays the accumulated statistics for the entire time interval.
<code>fsadm</code>	The <code>fsadm</code> command has the <code>-S</code> option to report shared block usage in the file system. You can use this option to find out the storage savings achieved through FileSnaps and how much real storage is required if all of the files are full copies. See the <code>fsadm_vxfs(1M)</code> manual page.
<code>fsmmap</code>	The <code>fsmmap</code> command has the <code>-c</code> option to report the count of the total number of physical blocks consumed by a file, and how many of those blocks might not be private to a given file. See the <code>fsmmap(1)</code> manual page.
<code>mkfs</code>	Use the <code>mkfs</code> command to make a disk layout Version 9 file system by specifying <code>-o version=9</code> . VxFS internally maintains a list of delayed operations on shared extent references and the size of this list (<code>rcqsize</code>) defaults to a value that is a function of the file system size, but can be changed when the file system is made. See the <code>mkfs_vxfs(1M)</code> manual page.

Table 25-1 (continued)

Command	Functionality
<code>vxfilesnap</code>	Use the <code>vxfilesnap</code> command to create a snapshot of a file or set of files or files in a directory. You can also use the <code>vxfilesnap</code> command to restore a older version of the file to the current file. See the <code>vxfilesnap(1)</code> manual page.
<code>vxtunefs</code>	The <code>vxtunefs</code> command supports an option to enable lazy copy-on-write tuneable, <code>lazy_copyonwrite</code>, on the file system, for better performance. See the <code>vxtunefs(1M)</code> manual page.

Using FileSnaps to create point-in-time copies of files

The key to obtaining maximum performance with FileSnaps is to minimize the copy-on-write overhead. You can achieved this by enabling lazy copy-on-write. Lazy copy-on-write is easy to enable and usually results in significantly better performance. If lazy copy-on-write is not a viable option for the use case under consideration, an efficient allocation of the source file can reduce the need of copy-on-write.

Using FileSnaps to provision virtual desktops

Virtual desktop infrastructure (VDI) operating system boot images are a good use case for FileSnaps. The parts of the boot images that can change are user profile, page files (or swap for UNIX/Linux) and application data. You should separate such data from boot images to minimize unsharing. You should allocate a single extent to the master boot image file.

Example of using FileSnaps to provision a virtual desktop

The following example uses a 4 GB master boot image that has a single extent that will be shared by all snapshots.

```
# touch /vdi_images/master_image
# /opt/VRTS/bin/setext -r 4g -f chgsize /vdi_images/master_image
```

The `master_image` file can be presented as a disk device to the virtual machine for installing the operating system. Once the operating system is installed and configured, the file is ready for snapshots.

Using FileSnaps to optimize write intensive applications for virtual machines

When virtual machines are spawned to perform certain tasks that are write intensive, a significant amount of unsharing can take place. Symantec recommends that you optimize performance by enabling lazy copy-on-write. If the use case does not allow enabling lazy copy-on-write, with careful planning, you can reduce the occurrence of unsharing. The easiest way to reduce unsharing is to separate the application data to a file other than the boot image. If you cannot do this due to the nature of your applications, then you can take actions similar to the following example.

Example of using FileSnaps to optimize write intensive applications

Assume that the disk space required for a boot image and the application data is 20 GB. Out of this, only 4 GB is used by the operating system and the remaining 16 GB is the space for applications to write. Any data or binaries that are required by each instance of the virtual machine can still be part of the first 4 GB of the shared extent. Since most of the writes are expected to take place on the 16 GB portion, you should allocate the master image in such a way that the 16 GB of space is not shared, as shown in the following commands:

```
# touch /vdi_images/master_image
# /opt/VRTS/bin/setext -r 4g -f chgsize /vdi_images/master_image
# dd if=/dev/zero of=/vdi_images/master_image seek=16777215 \
bs=1024 count=1
```

The last command creates a 16 GB hole at the end of the file. Since holes do not have any extents allocated, the writes to hole do not need to be unshared.

Using FileSnaps to create multiple copies of data instantly

It is common to create one or more copies of production data for the purpose of generating reports, mining, and testing. These cases frequently update the copies of the data with the most current data, and one or more copies of the data always exists. FileSnaps can be used to create multiple copies instantly. The application that uses the original data can see a slight performance hit due to the unsharing of data that can take place during updates. This slight impact on performance can still be present even when all FileSnaps have been deleted. However, you rarely see all FileSnaps being deleted since these use cases usually have one or more copies at any given time.

Comparison of the logical size output of the fsadm -S shared, du, and df commands

The `fsadm -S shared`, `du`, and `df` commands report different values for the size of a FileSnap. The `fsadm -S shared` command displays this size as the "logical size," which is the logical space consumed, in kilobytes, and accounts for both exclusive blocks and shared blocks. This value represents the actual disk space needed if the file system did not have any shared blocks. The value from the `fsadm -S shared` command differs from the output of `du -sk` command since the `du` command does not track the blocks consumed by VxFS structural files. As a result, the output of the `du -sk` command is less than the logical size output reported by the `fsadm -S shared` command.

The following examples show output from the `fsadm -S shared`, `du`, and `df` commands:

```
# mkfs -F vxfs -o version=9 /dev/vx/rdisk/dg/vol3
version 9 layout
104857600 sectors, 52428800 blocks of size 1024, log size 65536 blocks
rcq size 4096 blocks
largefiles supported

# mount -F vxfs /dev/vx/dsk/dg/vol3 /mnt

# df -k /mnt
Filesystem            1K-blocks    Used    Available Use% Mounted on
/dev/vx/dsk/dg1/vol3 52428800     83590   49073642   1%  /mnt

# /opt/VRTS/bin/fsadm -S shared /mnt
Mountpoint    Size(KB)    Available(KB)    Used(KB)    Logical_Size(KB)    Shared
/mnt          52428800    49073642         83590       83590               0%

# du -sk /mnt
0            /mnt

# dd if=/dev/zero of=/mnt/foo bs=1024 count=10
10+0 records in
10+0 records out
10240 bytes (10 kB) copied, 0.018901 seconds, 542 kB/s

# vxfilesnap /mnt/foo /mnt/foo.snap

# df -k /mnt
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/vx/dsk/dg1/vol3	52428800	83600	49073632	1%	/mnt

/opt/VRTS/bin/fsadm -S shared /mnt

Mountpoint	Size (KB)	Available (KB)	Used (KB)	Logical_Size (KB)	Shared
/mnt	52428800	49073632	83600	83610	0%

du -sk /mnt

20 /mnt

Administering snapshot file systems

This chapter includes the following topics:

- [Snapshot file system backups](#)
- [Snapshot file system performance](#)
- [About snapshot file system disk structure](#)
- [Differences between snapshots and Storage Checkpoints](#)
- [Creating a snapshot file system](#)
- [Examples of creating snapshot file systems](#)

Snapshot file system backups

After a snapshot file system is created, the snapshot maintains a consistent backup of data in the snapped file system.

Backup programs, such as `cpio`, that back up a standard file system tree can be used without modification on a snapshot file system because the snapshot presents the same data as the snapped file system. Backup programs, such as `vxdump`, that access the disk structures of a file system require some modifications to handle a snapshot file system.

VxFS utilities recognize snapshot file systems and modify their behavior so that they operate the same way on snapshots as they do on standard file systems. Other backup programs that typically read the raw disk image cannot work on snapshots without altering the backup procedure.

These other backup programs can use the `fscat` command to obtain a raw image of the entire file system that is identical to an image obtainable by running a `dd` command on the disk device containing the snapped file system at the exact moment the snapshot was created. The `snapread ioctl` takes arguments similar to those of the `read` system call and returns the same results that are obtainable by performing a read on the disk device containing the snapped file system at the exact time the snapshot was created. In both cases, however, the snapshot file system provides a consistent image of the snapped file system with all activity complete—it is an instantaneous read of the entire file system. This is much different than the results that would be obtained by a `dd` or `read` command on the disk device of an active file system.

Snapshot file system performance

Snapshot file systems maximize the performance of the snapshot at the expense of writes to the snapped file system. Reads from a snapshot file system typically perform at nearly the throughput rates of reads from a standard VxFS file system.

The performance of reads from the snapped file system are generally not affected. However, writes to the snapped file system, typically average two to three times as long as without a snapshot. This is because the initial write to a data block requires reading the old data, writing the data to the snapshot, and then writing the new data to the snapped file system. If there are multiple snapshots of the same snapped file system, writes are even slower. Only the initial write to a block experiences this delay, so operations such as writes to the intent log or inode updates proceed at normal speed after the initial write.

Reads from the snapshot file system are impacted if the snapped file system is busy because the snapshot reads are slowed by the disk I/O associated with the snapped file system.

The overall impact of the snapshot is dependent on the read to write ratio of an application and the mixing of the I/O operations. For example, a database application running an online transaction processing (OLTP) workload on a snapped file system was measured at about 15 to 20 percent slower than a file system that was not snapped.

About snapshot file system disk structure


A snapshot file system consists of:

- A super-block
- A bitmap

- A blockmap
- Data blocks copied from the snapped file system

The following figure shows the disk structure of a snapshot file system.

Figure 26-1 The Snapshot Disk Structure



The super-block is similar to the super-block of a standard VxFS file system, but the magic number is different and many of the fields are not applicable.

The bitmap contains one bit for every block on the snapped file system. Initially, all bitmap entries are zero. A set bit indicates that the appropriate block was copied from the snapped file system to the snapshot. In this case, the appropriate position in the blockmap references the copied block.

The blockmap contains one entry for each block on the snapped file system. Initially, all entries are zero. When a block is copied from the snapped file system to the snapshot, the appropriate entry in the blockmap is changed to contain the block number on the snapshot file system that holds the data from the snapped file system.

The data blocks are filled by data copied from the snapped file system, starting from the beginning of the data block area.

Differences between snapshots and Storage Checkpoints

While snapshots and Storage Checkpoints both create a point-in-time image of a file system and only the changed data blocks are updated, there are significant differences between the two technologies:

Table 26-1 Differences between snapshots and Storage Checkpoints

Snapshots	Storage Checkpoints
Require a separate device for storage	Reside on the same device as the original file system
Are read-only	Can be read-only or read-write
Are transient	Are persistent
Cease to exist after being unmounted	Can exist and be mounted on their own
Track changed blocks on the file system level	Track changed blocks on each file in the file system

Storage Checkpoints also serve as the enabling technology for two other Veritas features: Block-Level Incremental Backups and Storage Rollback, which are used extensively for backing up databases.

See [“About Storage Checkpoints”](#) on page 640.

Creating a snapshot file system

You create a snapshot file system by using the `-o snapof=` option of the `mount` command. The `-o snapsize=` option may also be required if the device you are mounting does not identify the device size in its disk label, or if you want a size smaller than the entire device.

You must make the snapshot file system large enough to hold any blocks on the snapped file system that may be written to while the snapshot file system exists. If a snapshot runs out of blocks to hold copied data, the snapshot is disabled and further attempts to access the snapshot file system fail.

During periods of low activity (such as nights and weekends), a snapshot typically requires about two to six percent of the blocks of the snapped file system. During a period of high activity, the snapshot of a typical file system may require 15 percent of the blocks of the snapped file system. Most file systems do not turn over 15 percent of data in a single day. These approximate percentages tend to be lower for larger file systems and higher for smaller file systems. You can allocate blocks to a snapshot based on characteristics such as file system usage and duration of backups.

Warning: Any existing data on the device used for the snapshot is overwritten.

To create a snapshot file system

- ◆ Mount the file system with the `-o snapof=` option:

```
# mount -F vxfs -o snapof=special,snapsize=snapshot_size \
  snapshot_special snapshot_mount_point
```

Examples of creating snapshot file systems

In the following examples, the `vxdump` utility is used to ascertain whether `/dev/vx/dsk/fsvol/vol1` is a snapshot mounted as `/backup/home` and does the appropriate work to get the snapshot data through the mount point.

These are typical examples of making a backup of a 300,000 block file system named `/home` using a snapshot file system on `/dev/vx/dsk/fsvol/vol1` with a snapshot mount point of `/backup/home`.

To create a backup using a snapshot file system

- 1 To back up files changed within the last week using `cpio`:

```
# mount -F vxfs -o snapof=/home,snapsize=100000 \
  /dev/vx/dsk/fsvol/vol1 /backup/home
# cd /backup
# find home -ctime -7 -depth -print | cpio -oc > \
  /dev/rmt/c0s0
# umount /backup/home
```

- 2 To do a level 3 backup of `/dev/vx/dsk/fsvol/vol1` and collect those files that have changed in the current directory:

```
# vxdump 3f - /dev/vx/dsk/fsvol/vol1 | vxrestore -xf -
```

- 3 To do a full backup of `/home`, which exists on disk `/dev/vx/dsk/fsvol/vol1`, and use `dd` to control blocking of output onto tape device using `vxdump`:

```
# mount -F vxfs -o snapof=/home,snapsize=100000 \
  /dev/vx/dsk/fsvol/vol1 /backup/home
# vxdump f - /dev/vx/dsk/fsvol/vol1 | dd bs=128k > \
  /dev/rmt/c0s0
```

Optimizing thin storage with Storage Foundation Cluster File System High Availability

- [Chapter 27. Understanding thin storage solutions in Storage Foundation Cluster File System High Availability](#)
- [Chapter 28. Migrating data from thick storage to thin storage](#)
- [Chapter 29. Maintaining Thin Storage with Thin Reclamation](#)

Understanding thin storage solutions in Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [About thin provisioning](#)
- [About thin optimization solutions in Storage Foundation Cluster File System High Availability](#)
- [About SmartMove](#)
- [About the Thin Reclamation feature](#)
- [Determining when to reclaim space on a thin reclamation LUN](#)
- [How automatic reclamation works](#)

About thin provisioning

See “[About SmartMove](#)” on page 675.

Thin provisioning is a storage array feature that optimizes storage use by allocating and reclaiming the storage on demand. With thin provisioning, the array allocates storage to applications only when the storage is needed, from a pool of free storage. Thin provisioning solves the problem of under-utilization of available array capacity. Administrators do not have to estimate how much storage an application requires. Instead, thin provisioning lets administrators provision large thin or thin reclaim

capable LUNs to a host. When the application writes data, the physical storage is allocated from the free pool on the array to the thin-provisioned LUNs.

The two types of thin provisioned LUNs are thin-capable or thin-reclaim capable. Both types of LUNs provide the capability to allocate storage as needed from the free pool. For example, storage is allocated when a file system creates or changes a file. However, this storage is not released to the free pool when files get deleted. Therefore, thin-provisioned LUNs can become 'thick' over time, as the file system starts to include unused free space where the data was deleted. Thin-reclaim capable LUNs address this problem with the ability to release the once-used storage to the pool of free storage. This operation is called thin storage reclamation.

The thin-reclaim capable LUNs do not perform the reclamation automatically. The server using the LUNs must initiate the reclamation. The administrator can initiate a reclamation manually, or with a scheduled reclamation operation.

Storage Foundation Cluster File System High Availability provides several features to support thin provisioning and thin reclamation, and to optimize storage use on thin provisioned arrays.

About thin optimization solutions in Storage Foundation Cluster File System High Availability

Array-based options like Thin Storage and Thin Provisioning help storage administrators to meet the challenges in managing their storage. These challenges include provisioning the storage, migrating data to maximize storage utilization, and maintaining the optimum storage utilization. Several features of Storage Foundation Cluster File System High Availability work together with the array functionality to solve these challenges.

[Table 27-1](#) lists the Storage Foundation Cluster File System High Availability features and benefits relating to thin storage.

Table 27-1 Thin storage solutions in Storage Foundation Cluster File System High Availability

Feature	Description	Benefits
SmartMove	The SmartMove feature moves or copies only blocks in use by the Veritas File System	Maximizes use of thin storage. See “About SmartMove” on page 675. Improves performance for copy operations. Enables migration from thick LUNs to thin provisioned LUNs. See “Migrating to thin provisioning” on page 679.
Thin disk discovery	Storage Foundation Cluster File System High Availability provides discovery for thin storage devices.	Recognizes and displays thin attributes for thin disks.
Thin Reclamation	Thin reclamation commands enable you to reclaim space on a file system, disk, disk group, or enclosure level.	Improves storage utilization and savings. See “About the Thin Reclamation feature” on page 676.

About SmartMove

Storage Foundation Cluster File System High Availability provides the SmartMove utility to optimize move and copy operations. The SmartMove utility leverages the knowledge that Veritas File System (VxFS) has of the Veritas Volume Manager (VxVM) storage. VxFS lets VxVM know which blocks have data. When VxVM performs an operation that copies or moves data, SmartMove enables the operation to only copy or move the blocks used by the file system. This capability improves performance for synchronization, mirroring, and copying operations because it reduces the number of blocks that are copied. SmartMove only works with VxFS file systems that are mounted on VxVM volumes. If a file system is not mounted, the utility has no visibility into the usage on the file system.

SmartMove is not used for volumes that have instant snapshots.

The SmartMove operation also can be used to migrate data from thick storage to thin-provisioned storage. Because SmartMove copies only blocks that are in use by the file system, the migration process results in a thin-provisioned LUN.

SmartMove for thin provisioning

Storage Foundation Cluster File System High Availability uses the SmartMove feature for thin provisioning. SmartMove enables you to migrate storage from thick storage to thin storage. SmartMove provides the ability to maintain the intent of thin provisioning.

Without SmartMove, synchronization between disks copies the entire storage that is allocated to Veritas File System (VxFS) and Veritas Volume Manager (VxVM). Synchronizing or resynchronizing a volume, plex, or subdisk can lead to unused space being allocated on the thin disk. Over time, normal operations cause the storage to become thick. With SmartMove, the disk synchronization copies only blocks that are actually in use at the file system level. This behavior prevents unused space from being allocated when a disk is synchronized or resynchronized. The disks stay thin.

The SmartMove feature is enabled for all disks by default. To take advantage of thin provisioning, SmartMove must be enabled at least for thin disks.

About the Thin Reclamation feature

Veritas Storage Foundation Cluster File System High Availability supports reclamation of the unused storage on thin-reclamation capable arrays. Storage Foundation Cluster File System High Availability automatically discovers LUNs that support thin reclamation.

A Veritas File System (VxFS) file system can be mounted on a Veritas Volume Manager (VxVM) volume that is backed by a thin-capable array. The size of the VxVM volume is a virtual size, that is backed by the free storage pool. When files are created or changed, storage is physically allocated to the file system from the array. If the files on the file system are deleted or shrunk in size, the space is freed from the file system usage. However, the space is not removed from the physical allocation. Over time, the physical space allocated to the file system is greater than the actual space used by the file system. The thin LUN eventually becomes 'thick', as the physical space allocated nears the size of the LUN.

The Thin Reclamation feature provides the ability to release this unused space back to the thin pool. Storage Foundation Cluster File System High Availability uses the VxFS allocation tables to identify unused blocks. VxVM maps this information about unused blocks down to the disk, enabling VxVM to return those blocks to the free pool. If the VxFS file system is not mounted, VxVM has no visibility into the file

system usage. Therefore, it is critical that the file system is mounted when you perform a reclamation. The operation of reclamation can be done on a disk group, LUN, enclosure, or file system.

VxVM reclaims space automatically when you delete a volume or remove a plex. The automatic reclamation is asynchronous, so that the space is not reclaimed at the array level immediately. The disk is marked as pending reclamation. You cannot remove a disk from VxVM until the reclamation completes. You can control the timing and frequency of the automatic reclamation.

Determining when to reclaim space on a thin reclamation LUN

When a thin LUN is used as a Veritas Volume Manager disk, the space is allocated only on an application write. Storage space is allocated from the free pool when files are created and written to in the file system. However, this storage is not automatically released to the free pool when data is deleted from a file system. As a result, all thin LUNs have a tendency to become thicker over time, with increased amounts of wasted storage (storage that is allocated but does not support application data).

As a storage administrator, you need to determine when to trigger the thin reclamation. The thin reclamation process can be time consuming, depending on various factors such as the size and fragmentation of the file system. The decision is a balance between how much space can be reclaimed, and how much time the reclaim operation will take.

The following considerations may apply:

- For a VxFS file system mounted on a VxVM volume, compare the file system usage to the actual physical allocation size to determine if a reclamation is desirable. If the file system usage is much smaller than the physical allocation size, it indicates that a lot of space can potentially be reclaimed. You may want to trigger a file system reclamation. If the file system usage is close to the physical allocation size, it indicates that the physical allocation is being used well. You may not want to trigger a reclamation.
See [“Displaying VxFS file system usage on thin reclamation LUNs”](#) on page 686.
- The array may provide notification when the storage pool usage has reached a certain threshold. You can evaluate whether you can reclaim space with Storage Foundation Cluster File System High Availability to free more space in the storage pool.
- Deleted volumes are reclaimed automatically. You can customize the schedule for automatic reclamation.

See [“Configuring automatic reclamation”](#) on page 692.

How automatic reclamation works

On thin-reclaimable arrays, storage that is no longer in use needs to be reclaimed by the array. Storage Foundation Cluster File System High Availability automatically reclaims the space on the array for certain administrative operations, as follows:

- Deleting a volume.
- Removing a mirror.
- Shrinking a volume.
- Removing a log.
- Creating or growing a volume with the `init=zero` option.

The process of reclaiming storage on an array can be intense on the array. To avoid any effect on regular I/O's to the array, Storage Foundation Cluster File System High Availability performs the reclaim operation asynchronously. The disk is flagged as pending reclamation. The `vxrelocd` (or recovery) daemon asynchronously reclaims the disks marked for reclamation at a future time. By default, the `vxrelocd` daemon runs every day at 22:10 hours, and reclaims storage on the deleted volumes or plexes that are one day old.

To display the disks that are pending reclamation, use the following command:

```
# vxprint -z
```

You can configure the automatic reclamation to reclaim immediately, or to schedule the asynchronous reclamation.

See [“Configuring automatic reclamation”](#) on page 692.

You can also trigger a reclamation manually for a disk, disk group or enclosure. This operation also reclaims any disks flagged as pending reclamation.

See [“Reclaiming space on a disk, disk group, or enclosure”](#) on page 689.

Migrating data from thick storage to thin storage

This chapter includes the following topics:

- [About using SmartMove to migrate to Thin Storage](#)
- [Migrating to thin provisioning](#)

About using SmartMove to migrate to Thin Storage

If you have existing data on a thick LUN, the SmartMove feature enables you to migrate the data to a thin LUN. The migration process copies only the blocks in use by the Veritas File System (VxFS) to the thin LUN. The SmartMove feature leverages the Veritas File System (VxFS) information about which blocks in a Veritas Volume Manager (VxVM) volume contain data. Therefore, the migration functionality is available only when a VxVM volume is on a mounted VxFS file system.

To migrate the data to the thin LUN, follow the recommended procedure.

See [“Migrating to thin provisioning”](#) on page 679.

Migrating to thin provisioning

The SmartMove™ feature enables migration from traditional LUNs to thinly provisioned LUNs, removing unused space in the process.

To migrate to thin provisioning

- 1 Check if the SmartMove feature is enabled.

```
# vxdefault list
KEYWORD                                CURRENT-VALUE  DEFAULT-VALUE
usefssmartmove                         all            all
...
```

If the output shows that the current value is none, configure SmartMove for all disks or thin disks.

See [“Configuring SmartMove”](#) on page 841.

- 2 Add the new, thin LUNs to the existing disk group. Enter the following commands:

```
# vxdisksetup -i da_name
# vxdg -g datadg adddisk da_name
```

where *da_name* is the disk access name in VxVM.

- 3 To identify LUNs with the `thinonly` or `thinrlm` attributes, enter:

```
# vxdisk -o thin list
```

- 4 Add the new, thin LUNs as a new plex to the volume.

NOTE: The VxFS file system must be mounted to get the benefits of the SmartMove feature.

The following methods are available to add the LUNs:

- Use the default settings for the `vxassist` command:

```
# vxassist -g datadg mirror datavol da_name
```

- Specify the `vxassist` command options for faster completion. The `-b` option copies blocks in the background. The following command improves I/O throughput:

```
# vxassist -b -oiosize=1m -t thinmig -g datadg mirror \
    datavol da_name
```

To view the status of the command, use the `vxtask` command:

```
# vxtask list
TASKID PTID TYPE/STATE    PCT    PROGRESS
211      ATCOPY/R 10.64% 0/20971520/2232320 PLXATT vol1 vol1-02 xivdg smartmove
```

```
212          ATCOPY/R 09.88% 0/20971520/2072576 PLXATT voll voll-03 xivdg smartmove
219          ATCOPY/R 00.27% 0/20971520/57344 PLXATT voll voll-04 xivdg smartmove
```

vxtask monitor 211

```
TASKID PTID TYPE/STATE    PCT    PROGRESS
211          ATCOPY/R 50.00% 0/20971520/10485760 PLXATT voll voll-02 xivdg smartmove
211          ATCOPY/R 50.02% 0/20971520/10489856 PLXATT voll voll-02 xivdg smartmove
211          ATCOPY/R 50.04% 0/20971520/10493952 PLXATT voll voll-02 xivdg smartmove
211          ATCOPY/R 50.06% 0/20971520/10498048 PLXATT voll voll-02 xivdg smartmove
211          ATCOPY/R 50.08% 0/20971520/10502144 PLXATT voll voll-02 xivdg smartmove
211          ATCOPY/R 50.10% 0/20971520/10506240 PLXATT voll voll-02 xivdg smartmove
```

- Specify the `vxassist` command options to reduce the effect on system performance. The following command takes longer to complete:

```
# vxassist -oslow -g datadg mirror datavol da_name
```

- 5 Optionally, test the performance of the new LUNs before removing the old LUNs.

To test the performance, use the following steps:

- Determine which plex corresponds to the thin LUNs:

```
# vxprint -g datadg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	datadg	datadg	-	-	-	-	-	-
dm	THINARRAY0_02	THINARRAY0_02	-	83886080	-	-	-	-
dm	STDARRAY1_01	STDARRAY1_01	-	41943040	-	-OHOTUSE	-	-
v	datavol	fsgen	ENABLED	41943040	-	ACTIVE	-	-
pl	datavol-01	datavol	ENABLED	41943040	-	ACTIVE	-	-
sd	STDARRAY1_01-01	datavol-01	ENABLED	41943040	0	-	-	-
pl	datavol-02	datavol	ENABLED	41943040	-	ACTIVE	-	-
sd	THINARRAY0_02-01	datavol-02	ENABLED	41943040	0	-	-	-

The example output indicates that the thin LUN corresponds to plex `datavol-02`.

- Direct all reads to come from those LUNs:

```
# vxvol -g datadg rdpol prefer datavol datavol-02
```

6 Remove the original non-thin LUNs.

Note: The ! character is a special character in some shells. This example shows how to escape it in a bash shell.

```
# vxassist -g datadg remove mirror datavol \!STDARRAY1_01
# vxdg -g datadg rmdisk STDARRAY1_01
# vxdisk rm STDARRAY1_01
```

7 Grow the file system and volume to use all of the larger thin LUN:

```
# vxresize -g datadg -x datavol 40g da_name
```

Maintaining Thin Storage with Thin Reclamation

This chapter includes the following topics:

- [Reclamation of storage on thin reclamation arrays](#)
- [Identifying thin and thin reclamation LUNs](#)
- [Displaying VxFS file system usage on thin reclamation LUNs](#)
- [Reclaiming space on a file system](#)
- [Reclaiming space on a disk, disk group, or enclosure](#)
- [Monitoring Thin Reclamation using the `vxtask` command](#)
- [Configuring automatic reclamation](#)

Reclamation of storage on thin reclamation arrays

Veritas Storage Foundation Cluster File System High Availability supports reclamation of the unused storage on thin-reclamation capable arrays and LUNs. Storage Foundation Cluster File System High Availability can reclaim blocks in a Veritas File System (VxFS) file system that is mounted on a Veritas Volume Manager (VxVM) volume.

The thin reclamation feature is supported only for LUNs that have the `thinrcldm` attribute. VxVM automatically discovers LUNs that support thin reclamation from thin capable storage arrays. You can list devices that are known to have the `thin` or `thinrcldm` attributes on the host.

See “[Identifying thin and thin reclamation LUNs](#)” on page 685.

For a list of the storage arrays that support thin reclamation, see the Symantec Hardware Compatibility List (HCL):

<http://www.symantec.com/docs/TECH170013>

Thin reclamation is not supported for boot devices.

You can use the thin reclamation feature in the following ways:

- Space is reclaimed automatically when a volume is deleted. Because it is asynchronous, you may not see the reclaimed space immediately.
- Perform the reclamation operation on a disk group, LUN, or enclosure using the `vxdisk` command.
See “[Reclaiming space on a disk, disk group, or enclosure](#)” on page 689.
- Perform the reclamation operation on a Veritas File System (VxFS) file system using the `fsadm` command.
See “[Reclaiming space on a file system](#)” on page 689.

Thin Reclamation of a disk, a disk group, or an enclosure

Storage Foundation Cluster File System High Availability provides the ability to reclaim unused space on thin-provisioned arrays, without needing to stop application I/O. The Veritas File System (VxFS) file system must be mounted.

You can trigger thin reclamation on one or more disks, disk groups, or enclosures. The reclamation process scans the specified storage for the VxVM volumes that have a mounted VxFS file system. Each volume is analyzed for any previously allocated space that the VxFS file system no longer uses. The unused space is released to the free storage pool on the thin array. The reclamation skips any volumes that do not have a mounted VxFS file system. The reclamation process also releases the space for any volumes or plexes that are marked as pending reclamation.

A full reclamation process also scans the specified storage for free space that is outside of the VxVM volumes.

Thin Reclamation takes a considerable amount of time when you reclaim thin storage on a large number of LUNs or an enclosure or disk group. As with other long-running operations, VxVM creates a task for a reclaim operation. You can monitor the reclaim operation with the `vxtask` command.

See “[Monitoring Thin Reclamation using the vxtask command](#)” on page 691.

Thin Reclamation of a file system

Veritas File System (VxFS) supports reclamation of free storage on a Thin Storage LUN. Free storage is reclaimed using the `fsadm` command or the `vxfs_ts_reclaim` API. You can perform the default reclamation or aggressive reclamation. If you used a file system for a long time and must perform reclamation on the file system, Symantec recommends that you run aggressive reclamation. Aggressive reclamation compacts the allocated blocks, which creates larger free blocks that can potentially be reclaimed.

See the `fsadm_vxfs(1M)` and `vxfs_ts_reclaim(3)` manual pages.

Thin Reclamation is only supported on file systems mounted on a VxVM volume.

Veritas File System also supports reclamation of a portion of the file system using the `vxfs_ts_reclaim()` API.

See the *Veritas File System Programmer's Reference Guide*.

Note: Thin Reclamation is a slow process and may take several hours to complete, depending on the file system size. Thin Reclamation is not guaranteed to reclaim 100% of the free space.

You can track the progress of the Thin Reclamation process by using the `vxtask list` command when using the Veritas Volume Manager (VxVM) command `vxdisk reclaim`.

See the `vxtask(1M)` and `vxdisk(1M)` manual pages.

You can administer Thin Reclamation using VxVM commands.

Identifying thin and thin reclamation LUNs

Using Veritas Dynamic Multi-Pathing (DMP), Storage Foundation Cluster File System High Availability automatically discovers thin devices that have been recognized on the host as `thin` or `thinreclm`. DMP uses the Veritas array support libraries (ASLs) to recognize vendor-specific thin attributes and claim devices accordingly as `thin` or `thinreclm`.

Thin devices that are classified as `thin` are capable of thin provisioning. Veritas Thin Reclamation only works on devices with the `thinreclm` attribute set. Before performing thin reclamation, determine whether the system recognizes the LUN as a `thinreclm` host.

To identify devices on a host that are known to have the `thin` or `thinreclm` attributes, use the `vxdisk -o thin list` command. The `vxdisk -o thin list` command

also reports on the size of the disk, and the physical space that is allocated on the array.

To identify thin and thinrclm LUNs

- ◆ To identify all of the `thin` or `thinrclm` LUNs that are known to a host, use the following command:

```
# vxdisk -o thin list
```

DEVICE	SIZE (mb)	PHYS_ALLOC (mb)	GROUP	TYPE
hitachi_usp0_065a	10000	84	-	thinrclm
hitachi_usp0_065b	10000	110	-	thinrclm
hitachi_usp0_065c	10000	74	-	thinrclm
hitachi_usp0_065d	10000	50	-	thinrclm
.				
.				
.				
hitachi_usp0_0660	10000	672	thindg	thinrclm
emc_clariion0_48	30720	N/A	-	thin
emc_clariion0_49	30720	N/A	-	thin
emc_clariion0_50	30720	N/A	-	thin
emc_clariion0_51	30720	N/A	-	thin

In the output, the `SIZE` column shows the size of the disk. The `PHYS_ALLOC` column shows the physical allocation on the array side. The `TYPE` indicates whether the array is thin or thinrclm.

See the `vxdisk(1m)` manual page.

Displaying VxFS file system usage on thin reclamation LUNs

Storage Foundation Cluster File System High Availability can discover and display the disk space usage for Veritas File System (VxFS) file systems on `thin` or `thinrclm` devices. The VxFS file systems must be mounted on Veritas Volume Manager (VxVM) volumes. The usage information can help you decide when to perform thin reclamation of a file system.

See [“Determining when to reclaim space on a thin reclamation LUN”](#) on page 677.

To report the disk space usage for mounted VxFS file systems on VxVM volumes, use the `vxdisk -o thin -o fssize list` command. The command displays the amount of disk space that currently contains files and is actively in use by the VxFS file system. The usage does not include any space that is allocated to the file system

but was freed by deleting files. If more than one mounted VxFS file system uses the device, the file system usage column displays the consolidated space usage.

The following limitations apply to the command to display file system usage:

- The `-o fssize` option does not display the space used by cache objects or instant snapshots.
- RAID5 format is not supported.
- If the VxFS file system is not mounted, or if the device has both mounted and unmounted VxFS file systems, no information is displayed. The file system usage (FS_USAGE) column displays a dash (-).

You can display the size and usage for all `thin` or `thinrclm` LUNs, or specify an enclosure name or a device name. If you specify one or more devices or enclosures, the command displays only the space usage on the specified devices. If the specified device is not a `thin` device or `thinrclm` device, the device is listed but the FS_USAGE column displays a dash (-).

If a VxFS file system spans multiple devices, you must specify all of the devices to display the entire file system usage. If you specify only some of the devices, the file system usage is incomplete. The command ignores the file system usage on any devices that are not specified.

Note: The command can potentially take a long time to complete depending on the file system size, the level of fragmentation, and other factors. The command creates a task that you can monitor with the `vxtask` command.

The command output displays the following information.

SIZE	The size of the disk; that is, the size that is presented to the file system. This size represents the virtual size rather than the actual physical space used on the device.
PHYS_ALLOC	The physical allocation on the array side. This size represents the physical space that is allocated as the application writes to the file system. When the files are deleted or changed, the physical space remains allocated until a reclamation is performed. In this case, the physical size includes some unused space.
FS_USAGE	The physical space Veritas File System (VxFS) file systems are using. The VxFS file systems must be mounted on VxVM volumes. The information is displayed only for thin provisioning capable (thin) or thin reclamation capable (thinrclm) LUNs.

GROUP	The disk group that contains the disk.
TYPE	The type of thin devices – thin provisioning capable (thin) or thin reclamation capable (thinrlm). The <code>vxdisk -o thin list</code> command displays thick disks only if you explicitly specify the disk name on the command line.

To display file system usage on all thin LUNs

- ◆ To display the file system usage on all `thin` or `thinrlm` LUNs known to the system, use the following command:

```
$ vxdisk -o thin,fssize [-u unit] list
```

Where *unit* is a size unit for the display. For example:

```
$ vxdisk -o thin,fssize -u m list
DEVICE      SIZE      PHYS_ALLOC FS_USAGE GROUP TYPE
emc0_428a  16384.00m 6335.00m   610.00m  mydg  thinrlm
emc0_428b  16384.00m 3200.00m    22.00m  mydg  thinrlm
emc0_4287  16384.00m 6233.00m   617.00m  mydg  thinrlm
emc0_4288  16384.00m 1584.00m   617.00m  mydg  thinrlm
emc0_4289  16384.00m 2844.00m  1187.00m  mydg  thinrlm
xiv0_030f  16384.00m 2839.00m  1223.00m  xivdg thinrlm
xiv0_0307  16384.00m  666.00m   146.00m  xivdg thinrlm
xiv0_0308  16384.00m  667.00m   147.00m  xivdg thinrlm
xiv0_0309  16384.00m   3.00m    -        -      thinrlm
xiv0_0310  16384.00m  30.00m    -        -      thinrlm
```

Or, to display the file system usage on a specific LUN or enclosure, use the following form of the command:

```
$ vxdisk -o thin,fssize list [-u unit] disk|enclosure
```

For example:

```
$ vxdisk -o thin,fssize list emc0
DEVICE      SIZE (mb)  PHYS_ALLOC (mb)  FS_USAGE (mb)  GROUP  TYPE
emc0_428a  16384     6335            610            mydg   thinrlm
emc0_428b  16384     6335            624            mydg   thinrlm
emc0_4287  16384     6335            617            mydg   thinrlm
emc0_4288  16384     1584            617            mydg   thinrlm
emc0_4289  16384     2844            1187           mydg   thinrlm
```

Reclaiming space on a file system

This section describes how to trigger space reclamation on a file system. You can perform a default space reclamation, or aggressive reclamation. Thin Reclamation is only supported on VxFS file systems mounted on a VxVM volume.

Thin Reclamation is not supported for file systems mounted on RAID5 volumes.

To trigger aggressive space reclamation

- 1 Ensure you mounted the VxFS file system.

See the `mount(1M)` manual page.

If you need to mount the VxFS file system, see the `mount_vxfs(1M)` manual page.

- 2 Use the following command to perform aggressive reclamation of free storage to the Thin Storage LUN on a VxFS file system mounted at `/mnt1`:

```
# /opt/VRTS/bin/fsadm -R -o aggressive /mnt1
```

To trigger space reclamation

- 1 Ensure you mounted the VxFS file system.

See the `mount(1M)` manual page.

If you need to mount the VxFS file system, see the `mount_vxfs(1M)` manual page.

- 2 Use the `fsadm` command to trigger space reclamation:

```
# /opt/VRTS/bin/fsadm -F vxfs -R /<VxFS_mount_point>
```

where `<VxFS_mount_point>` is the name of the VxFS file system mount point.

Note: If the VxFS file system is not mounted you will receive an error message. For example: Disk 3pardata0_110 : Skipped. No VxFS file system found.

Reclaiming space on a disk, disk group, or enclosure

Use the `vxdisk reclaim` command to trigger online Thin Reclamation on one or more disks, disk groups, or enclosures. By default, the `vxdisk reclaim` command performs Thin Reclamation on the disks where the VxVM volume is on a “mounted” VxFS file system. The reclamation skips disks that do not have a VxFS file system mounted. Thin reclamation is not supported for RAID-5 volumes, or for instant snapshots.

Reclaiming space on a disk

- ◆ Use the following command to trigger reclamation:

```
# vxdisk reclaim [disk...]
```

For example, to trigger reclamation on LUNs hitachi_usp0_065a and hitachi_usp0_065b:

```
# vxdisk reclaim hitachi_usp0_065a hitachi_usp0_065b
```

In the above example, suppose the hitachi_usp0_065a contains a VxVM volume vol1 with a VxFS file system. If the VxFS file system is not mounted, the command skips reclamation for hitachi_usp0_065a. The command scans hitachi_usp0_065b, and reclaims any unused space.

Performing an aggressive space reclamation on a disk

- ◆ Use the following command to trigger reclamation:

```
# vxdisk -o full reclaim [disk...]
```

For example, to trigger reclamation on LUNs hitachi_usp0_065a:

```
# vxdisk -o full reclaim hitachi_usp0_065a
```

In the above example, suppose the hitachi_usp0_065a contains a VxVM volume vol1 with a VxFS file system mounted. With the -o full option, the above command scans hitachi_usp0_065a for unused space outside of the vol1, and reclaims any unused space found. For example, if there is space between subdisks, it is reclaimed.

Reclaiming space on a disk group

- ◆ Use the following command to trigger reclamation:

```
# vxdisk reclaim diskgroup
```

For example, to trigger reclamation on the disk group oradg:

```
# vxdisk reclaim oradg
```

Reclaiming space on an enclosure

- ◆ Use the following command to trigger reclamation:

```
# vxdisk reclaim enclosure
```

For example, to trigger reclamation on the enclosure=EMC_CLARiiON0:

```
# vxdisk reclaim EMC_CLARiiON0
```

Monitoring Thin Reclamation using the `vxtask` command

The thin reclamation can be an intensive operation that may be time consuming, depending on the size of the disk and the amount of space to be reclaimed. As with other long-running tasks, you can monitor the operation with the `vxtask` command.

To monitor thin reclamation

- 1 Initiate the thin reclamation as usual, for a disk, disk group, or enclosure.

```
# vxdisk reclaim diskgroup| disk| enclosure
```

For example:

```
# vxdisk reclaim dg100
```

- 2 To monitor the reclamation status, run the following command in another session:

```
# vxtask monitor
```

TASKID	PTID	TYPE/STATE	PCT	PROGRESS
1258	-	RECLAIM/R	17.28%	65792/33447328/5834752 RECLAIM vol14 dg100
1259	-	RECLAIM/R	25.98%	0/20971520/5447680 RECLAIM vol12 dg100
1263	-	RECLAIM/R	25.21%	0/20971520/5287936 RECLAIM vol13 dg100
1258	-	RECLAIM/R	25.49%	0/20971520/3248128 RECLAIM vol14 dg100
1258	-	RECLAIM/R	27.51%	0/20971520/3252224 RECLAIM vol14 dg100
1263	-	RECLAIM/R	25.23%	0/20971520/5292032 RECLAIM vol13 dg100
1259	-	RECLAIM/R	26.00%	0/20971520/5451776 RECLAIM vol12 dg100

- 3 If you have multiple tasks, you can use the following command to display the tasks.

```
# vxtask list
```

TASKID	PTID	TYPE/STATE	PCT	PROGRESS
1258	-	RECLAIM/R	17.28%	65792/33447328/5834752 RECLAIM vol4 dg100
1259	-	RECLAIM/R	25.98%	0/20971520/5447680 RECLAIM vol2 dg100
1263	-	RECLAIM/R	25.21%	0/20971520/5287936 RECLAIM vol3 dg100

- 4 Use the task id from the above output to monitor the task:

```
# vxtask monitor 1258
```

TASKID	PTID	TYPE/STATE	PCT	PROGRESS
1258	-	RECLAIM/R	17.28%	65792/33447328/5834752 RECLAIM vol4 dg100
1258	-	RECLAIM/R	32.99%	65792/33447328/11077632 RECLAIM vol4 dg100
1258	-	RECLAIM/R	45.55%	65792/33447328/15271936 RECLAIM vol4 dg100
1258	-	RECLAIM/R	50.00%	0/20971520/10485760 RECLAIM vol4 dg100
.				
.				
.				

The `vxdisk reclaim` command runs in another session while you run the `vxtask list` command.

See the `vxtask(1m)` manual page.

Configuring automatic reclamation

The `vxrelocd` daemon tracks the disks that require reclamation. By default, the `vxrelocd` daemon runs everyday at 22:10 hours and reclaims storage on the deleted volume that are one day old.

To control the schedule for reclamation, use the following tunable parameters:

`reclaim_on_delete_wait_period`

Specifies the number of days after a volume or plex is deleted when VxVM reclaims the storage space. The value is an integer between -1 and 367.

The default value is 1, which means the volume is deleted the next day.

A value of -1 indicates that the storage is reclaimed immediately.

A value of 367 indicates that the storage space is not reclaimed automatically. Storage space can only be reclaimed manually using the `vxdisk reclaim` command.

`reclaim_on_delete_start_time`

The time of day when VxVM starts the reclamation for deleted volumes. The value is any time of day in 24 hour format. (hh:mm)

The default time is 22:10.

Change the tunables using the `vxdefault` command. See the `vxdefault(1m)` manual page.

Maximizing storage utilization

- [Chapter 30. Understanding storage tiering with SmartTier](#)
- [Chapter 31. Creating and administering volume sets](#)
- [Chapter 32. Multi-volume file systems](#)
- [Chapter 33. Administering SmartTier](#)
- [Chapter 34. Administering hot-relocation](#)
- [Chapter 35. Deduplicating data on Solaris SPARC](#)
- [Chapter 36. Compressing files](#)

Understanding storage tiering with SmartTier

This chapter includes the following topics:

- [About SmartTier](#)
- [SmartTier building blocks](#)
- [How SmartTier works](#)
- [SmartTier in a High Availability \(HA\) environment](#)

About SmartTier

SmartTier matches data storage with data usage requirements. After data matching, the data can then be relocated based upon data usage and other requirements determined by the storage or database administrator (DBA).

As more and more data is retained over a period of time, eventually, some of that data is needed less frequently. The data that is needed less frequently still requires a large amount of disk space. SmartTier enables the database administrator to manage data so that less frequently used data can be moved to slower, less expensive disks. This also permits the frequently accessed data to be stored on faster disks for quicker retrieval.

Tiered storage is the assignment of different types of data to different storage types to improve performance and reduce costs. With SmartTier, storage classes are used to designate which disks make up a particular tier. There are two common ways of defining storage classes:

Note: SmartTier is the expanded and renamed feature previously known as Dynamic Storage Tiering (DST).

- Performance, or storage, cost class: The most-used class consists of fast, expensive disks. When data is no longer needed on a regular basis, the data can be moved to a different class that is made up of slower, less expensive disks.
- Resilience class: Each class consists of non-mirrored volumes, mirrored volumes, and n-way mirrored volumes.
For example, a database is usually made up of data, an index, and logs. The data could be set up with a three-way mirror because data is critical. The index could be set up with a two-way mirror because the index is important, but can be recreated. The redo and archive logs are not required on a daily basis but are vital to database recovery and should also be mirrored.

SmartTier policies control initial file location and the circumstances under which existing files are relocated. These policies cause the files to which they apply to be created and extended on specific subsets of a file system's volume set, known as placement classes. The files are relocated to volumes in other placement classes when they meet specified naming, timing, access rate, and storage capacity-related conditions.

In addition to preset policies, you can manually move files to faster or slower storage with SmartTier, when necessary. You can also run reports that list active policies, display file activity, display volume usage, or show file statistics.

To use SmartTier, your storage must be managed using the following features:

- VxFS multi-volume file system
- VxVM volume set
- Volume tags
- SmartTier management at the file level
- SmartTier management at the sub-file level

SmartTier building blocks

To use SmartTier, your storage must be managed using the following features:

- VxFS multi-volume file system
- VxVM volume set
- Volume tags

- SmartTier management at the file level
- SmartTier management at the sub-file level

About VxFS multi-volume file systems

Multi-volume file systems are file systems that occupy two or more virtual volumes. The collection of volumes is known as a volume set, and is made up of disks or disk array LUNs belonging to a single Veritas Volume Manager (VxVM) disk group. A multi-volume file system presents a single name space, making the existence of multiple volumes transparent to users and applications. Each volume retains a separate identity for administrative purposes, making it possible to control the locations to which individual files are directed.

See [“About multi-volume file systems”](#) on page 709.

This feature is available only on file systems meeting the following requirements:

- The minimum disk group version is 140.
- The minimum file system layout version is 7 for file level SmartTier.
- The minimum file system layout version is 8 for sub-file level SmartTier.

To convert your existing VxFS system to a VxFS multi-volume file system, you must convert a single volume to a volume set.

The VxFS volume administration utility (fsvoladm utility) can be used to administer VxFS volumes. The fsvoladm utility performs administrative tasks, such as adding, removing, resizing, encapsulating volumes, and setting, clearing, or querying flags on volumes in a specified Veritas File System.

See the `fsvoladm` (1M) manual page for additional information about using this utility.

About VxVM volume sets

Volume sets allow several volumes to be represented by a single logical object. Volume sets cannot be empty. All I/O from and to the underlying volumes is directed via the I/O interfaces of the volume set. The volume set feature supports the multi-volume enhancement to Veritas File System (VxFS). This feature allows file systems to make best use of the different performance and availability characteristics of the underlying volumes. For example, file system metadata could be stored on volumes with higher redundancy, and user data on volumes with better performance.

About volume tags

You make a VxVM volume part of a placement class by associating a volume tag with it. For file placement purposes, VxFS treats all of the volumes in a placement class as equivalent, and balances space allocation across them. A volume may have more than one tag associated with it. If a volume has multiple tags, the volume belongs to multiple placement classes and is subject to allocation and relocation policies that relate to any of the placement classes.

Warning: Multiple tagging should be used carefully.

A placement class is a SmartTier attribute of a given volume in a volume set of a multi-volume file system. This attribute is a character string, and is known as a volume tag.

How SmartTier works

SmartTier is a VxFS feature that enables you to allocate file storage space from different storage tiers according to rules you create. SmartTier provides a more flexible alternative compared to current approaches for tiered storage. Static storage tiering involves a manual one-time assignment of application files to a storage class, which is inflexible over a long term. Hierarchical Storage Management solutions typically require files to be migrated back into a file system name space before an application access request can be fulfilled, leading to latency and run-time overhead. In contrast, SmartTier allows organizations to:

- Optimize storage assets by dynamically moving a file to its optimal storage tier as the value of the file changes over time
- Automate the movement of data between storage tiers without changing the way users or applications access the files
- Migrate data automatically based on policies set up by administrators, eliminating operational requirements for tiered storage and downtime commonly associated with data movement

SmartTier leverages two key technologies included with Veritas Storage Foundation Cluster File System High Availability: support for multi-volume file systems and automatic policy-based placement of files within the storage managed by a file system. A multi-volume file system occupies two or more virtual storage volumes and thereby enables a single file system to span across multiple, possibly heterogeneous, physical storage devices. For example the first volume could reside on EMC Symmetrix DMX spindles, and the second volume could reside on EMC CLARiiON spindles. By presenting a single name space, multi-volumes are

transparent to users and applications. This multi-volume file system remains aware of each volume's identity, making it possible to control the locations at which individual files are stored. When combined with the automatic policy-based placement of files, the multi-volume file system provides an ideal storage tiering facility, which moves data automatically without any downtime requirements for applications and users alike.

In a database environment, the access age rule can be applied to some files. However, some data files, for instance are updated every time they are accessed and hence access age rules cannot be used. SmartTier provides mechanisms to relocate portions of files as well as entire files to a secondary tier.

SmartTier file management

SmartTier enables administrators of multi-volume VxFS file systems to manage the placement of files on individual volumes in a volume set by defining placement policies that control both initial file location and the circumstances under which existing files are relocated. These placement policies cause the files to which they apply to be created and extended on specific subsets of a file system's volume set, known as placement classes. The files are relocated to volumes in other placement classes when they meet the specified naming, timing, access rate, and storage capacity-related conditions.

File-based movement:

- The administrator can create a file allocation policy based on filename extension before new files are created, which will create the datafiles on the appropriate tier during database creation.
- The administrator can also create a file relocation policy for database files or any types of files, which would relocate files based on how frequently a file is used.

SmartTier sub-file object management

SmartTier enables administrators of multi-volume VxFS file systems to manage the placement of file objects as well as entire files on individual volumes.

Using sub-file based movement you can:

- Move a set of ranges of a specified set of files of a specified set of mounts to a desired set of tiers on command.
- Move segments of files using automation to:
 - Monitor a set of files for collecting I/O statistics
 - Periodically collect and persist the statistics, cluster-wide if applicable

- Periodically enforce the ranges of the registered sets of files based on their relative frequency of access to a desired set of tiers
- Track the historical movements of those ranges

How the SmartTier policy works with the shared extents

The SmartTier enforcement operation ignores moving the shared extents. For example, consider a file A that contains some shared and private extents that belong to device 1. If the user sets a policy that states that all the extents of the file A must be allocated to device 2, the SmartTier enforcement operation moves all the non-shared extents from device 1 to device 2. However, the SmartTier enforcement operation ignores moving the shared extents. As a result, the file A still contains shared extents that belong to device 1. This occurs even after the successful execution of the SmartTier enforcement operation.

On the other hand, any subsequent new allocation on behalf of the file A adheres to the preset SmartTier policy. Since the copy-on-write or unshare operation requires a new allocation, the SmartTier enforcement operation complies with the preset policy. If a write operation on the file A writes to shared extents, new allocations as part of copy-on-write operation is done from device 2. This behaviour adheres to the preset SmartTier policy.

SmartTier in a High Availability (HA) environment

Veritas Cluster Server does not provide a bundled agent for volume sets. If issues arise with volumes or volume sets, the issues can only be detected at the DiskGroup and Mount resource levels.

The DiskGroup agent brings online, takes offline, and monitors a Veritas Volume Manager (VxVM) disk group. This agent uses VxVM commands. When the value of the StartVolumes and StopVolumes attributes are both 1, the DiskGroup agent online and offlines the volumes during the import and deport operations of the disk group. When using volume sets, set StartVolumes and StopVolumes attributes of the DiskGroup resource that contains the volume are set to 1. If a file system is created on the volume set, use a Mount resource to mount the volume set.

The Mount agent brings online, takes offline, and monitors a file system or NFS client mount point.

For additional information, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

Creating and administering volume sets

This chapter includes the following topics:

- [About volume sets](#)
- [Creating a volume set](#)
- [Adding a volume to a volume set](#)
- [Removing a volume from a volume set](#)
- [Listing details of volume sets](#)
- [Stopping and starting volume sets](#)
- [Managing raw device nodes of component volumes](#)

About volume sets

Veritas File System (VxFS) uses volume sets to implement its Multi-Volume Support and SmartTier features.

See “[About SmartTier](#)” on page 695.

Veritas Volume Manager (VxVM) provides the `vxvset` command to create and administer volume sets.

See the `vxvset(1M)` manual page.

Volume sets have the following limitations:

- A maximum of 2048 volumes can be configured in a volume set.
- Only a Veritas File System is supported on a volume set.

- The first volume (index 0) in a volume set must be larger than the sum of the total volume size divided by 4000, the size of the VxFS intent log, and 1MB. Volumes 258 MB or larger should always suffice.
- Raw I/O from and to a volume set is not supported.
- Raw I/O from and to the component volumes of a volume set is supported under certain conditions.
See [“Managing raw device nodes of component volumes”](#) on page 705.
- Volume sets can be used in place of volumes with the following `vxsnap` operations on instant snapshots: `addmir`, `dis`, `make`, `prepare`, `reattach`, `refresh`, `restore`, `rmmir`, `split`, `syncpause`, `syncresume`, `syncstart`, `syncstop`, `syncwait`, and `unprepare`. The third-mirror break-off usage model for full-sized instant snapshots is supported for volume sets provided that sufficient plexes exist for each volume in the volume set.
For more information about snapshots, see the *Veritas Storage Foundation and High Availability Solutions Solutions Guide*.
- A full-sized snapshot of a volume set must itself be a volume set with the same number of volumes and the same volume index numbers as the parent. The corresponding volumes in the parent and snapshot volume sets are also subject to the same restrictions as apply between standalone volumes and their snapshots.

Creating a volume set

To create a volume set for use by Veritas File System (VxFS), use the following command:

```
# vxvset [-g diskgroup] -t vxfs make volset  
      volume
```

Here *volset* is the name of the volume set, and *volume* is the name of the first volume in the volume set. The `-t vxfs` option creates the volume set configured for use by VxFS. You must create the volume before running the command. `vxvset` will not automatically create the volume.

For example, to create a volume set named `myvset` that contains the volume `voll`, in the disk group `mydg`, you would use the following command:

```
# vxvset -g mydg -t vxfs make myvset voll
```

Adding a volume to a volume set

Having created a volume set containing a single volume, you can use the following command to add further volumes to the volume set:

```
# vxvset [-g diskgroup] [-f] addvol volset  
      volume
```

For example, to add the volume `vol2`, to the volume set `myvset`, use the following command:

```
# vxvset -g mydg addvol myvset vol2
```

Warning: The `-f` (force) option must be specified if the volume being added, or any volume in the volume set, is either a snapshot or the parent of a snapshot. Using this option can potentially cause inconsistencies in a snapshot hierarchy if any of the volumes involved in the operation is already in a snapshot chain.

Removing a volume from a volume set

To remove a component volume from a volume set, use the following command:

```
# vxvset [-g diskgroup] [-f] rmvol volset  
      volume
```

For example, the following commands remove the volumes, `vol1` and `vol2`, from the volume set `myvset`:

```
# vxvset -g mydg rmvol myvset vol1  
# vxvset -g mydg rmvol myvset vol2
```

Removing the final volume deletes the volume set.

Warning: The `-f` (force) option must be specified if the volume being removed, or any volume in the volume set, is either a snapshot or the parent of a snapshot. Using this option can potentially cause inconsistencies in a snapshot hierarchy if any of the volumes involved in the operation is already in a snapshot chain.

Listing details of volume sets

To list the details of the component volumes of a volume set, use the following command:

```
# vxvset [-g diskgroup] list [volset]
```

If the name of a volume set is not specified, the command lists the details of all volume sets in a disk group, as shown in the following example:

```
# vxvset -g mydg list
```

NAME	GROUP	NVOLS	CONTEXT
set1	mydg	3	-
set2	mydg	2	-

To list the details of each volume in a volume set, specify the name of the volume set as an argument to the command:

```
# vxvset -g mydg list set1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
vol1	0	12582912	ENABLED	-
vol2	1	12582912	ENABLED	-
vol3	2	12582912	ENABLED	-

The context field contains details of any string that the application has set up for the volume or volume set to tag its purpose.

Stopping and starting volume sets

Under some circumstances, you may need to stop and restart a volume set. For example, a volume within the set may have become detached, as shown here:

```
# vxvset -g mydg list set1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
vol1	0	12582912	DETACHED	-
vol2	1	12582912	ENABLED	-
vol3	2	12582912	ENABLED	-

To stop and restart one or more volume sets, use the following commands:

```
# vxvset [-g diskgroup] stop volset ...
# vxvset [-g diskgroup] start volset ...
```

For the example given previously, the effect of running these commands on the component volumes is shown below:

```
# vxvset -g mydg stop set1
```

```
# vxvset -g mydg list set1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
vol1	0	12582912	DISABLED	-
vol2	1	12582912	DISABLED	-
vol3	2	12582912	DISABLED	-

```
# vxvset -g mydg start set1
```

```
# vxvset -g mydg list set1
```

VOLUME	INDEX	LENGTH	KSTATE	CONTEXT
vol1	0	12582912	ENABLED	-
vol2	1	12582912	ENABLED	-
vol3	2	12582912	ENABLED	-

Managing raw device nodes of component volumes

To guard against accidental file system and data corruption, the device nodes of the component volumes are configured by default not to have raw and block entries in the `/dev/vx/rdisk/diskgroup` and `/dev/vx/dsk/diskgroup` directories. As a result, applications are prevented from directly reading from or writing to the component volumes of a volume set.

If some applications, such as the raw volume backup and restore feature of the Symantec NetBackup™ software, need to read from or write to the component volumes by accessing raw device nodes in the `/dev/vx/rdisk/diskgroup` directory, this is supported by specifying additional command-line options to the `vxvset` command. Access to the block device nodes of the component volumes of a volume set is unsupported.

Warning: Writing directly to or reading from the raw device node of a component volume of a volume set should only be performed if it is known that the volume's data will not otherwise change during the period of access.

All of the raw device nodes for the component volumes of a volume set can be created or removed in a single operation. Raw device nodes for any volumes added to a volume set are created automatically as required, and inherit the access mode of the existing device nodes.

Access to the raw device nodes for the component volumes can be configured to be read-only or read-write. This mode is shared by all the raw device nodes for the component volumes of a volume set. The read-only access mode implies that any writes to the raw device will fail, however writes using the `ioctl` interface or by VxFS to update metadata are not prevented. The read-write access mode allows direct writes via the raw device. The access mode to the raw device nodes of a volume set can be changed as required.

The presence of raw device nodes and their access mode is persistent across system reboots.

Note the following limitations of this feature:

- The disk group version must be 140 or greater.
- Access to the raw device nodes of the component volumes of a volume set is only supported for private disk groups; it is not supported for shared disk groups in a cluster.

Enabling raw device access when creating a volume set

To enable raw device access when creating a volume set, use the following form of the `vxvset make` command:

```
# vxvset [-g diskgroup] -o makedev=on \
  [-o compvol_access={read-only|read-write}] \
  [-o index] [-c "ch_addopt"] make vset
  vol [index]
```

The `-o makedev=on` option enables the creation of raw device nodes for the component volumes at the same time that the volume set is created. The default setting is `off`.

If the `-o compvol_access=read-write` option is specified, direct writes are allowed to the raw device of each component volume. If the value is set to `read-only`, only reads are allowed from the raw device of each component volume.

If the `-o makedev=on` option is specified, but `-o compvol_access` is not specified, the default access mode is `read-only`.

If the `vxvset addvol` command is subsequently used to add a volume to a volume set, a new raw device node is created in `/dev/vx/rdisk/diskgroup` if the value of the `makedev` attribute is currently set to `on`. The access mode is determined by the current setting of the `compvol_access` attribute.

The following example creates a volume set, `myvset1`, containing the volume, `myvol1`, in the disk group, `mydg`, with raw device access enabled in read-write mode:

```
# vxvset -g mydg -o makedev=on -o compvol_access=read-write \
    make myvset1 myvol1
```

Displaying the raw device access settings for a volume set

You can use the `vxprint -m` command to display the current settings for a volume set. If the `makedev` attribute is set to `on`, one of the following strings is displayed in the output:

```
vset_devinfo=on:read-only      Raw device nodes in read-only mode.
vset_devinfo=on:read-write     Raw device nodes in read-write mode.
```

A string is not displayed if `makedev` is set to `off`.

If the output from the `vxprint -m` command is fed to the `vxmlake` command to recreate a volume set, the `vset_devinfo` attribute must set to `off`. Use the `vxvset` set command to re-enable raw device access with the desired access mode.

See [“Controlling raw device access for an existing volume set”](#) on page 707.

Controlling raw device access for an existing volume set

To enable or disable raw device node access for an existing volume set, use the following command:

```
# vxvset [-g diskgroup] [-f] set makedev={on|off} vset
```

The `makedev` attribute can be specified to the `vxvset set` command to create (`makedev=on`) or remove (`makedev=off`) the raw device nodes for the component volumes of a volume set. If any of the component volumes are open, the `-f` (force) option must be specified to set the attribute to `off`.

Specifying `makedev=off` removes the existing raw device nodes from the `/dev/vx/rdisk/diskgroup` directory.

If the `makedev` attribute is set to `off`, and you use the `mknod` command to create the raw device nodes, you cannot read from or write to those nodes unless you set the value of `makedev` to `on`.

The syntax for setting the `compvol_access` attribute on a volume set is:

```
# vxvset [-g diskgroup] [-f] set \
    compvol_access={read-only|read-write} vset
```

The `compvol_access` attribute can be specified to the `vxvset set` command to change the access mode to the component volumes of a volume set. If any of the

component volumes are open, the `-f` (force) option must be specified to set the attribute to `read-only`.

The following example sets the `makedev=on` and `compvol_access=read-only` attributes on a volume set, `myvset2`, in the disk group, `mydg`:

```
# vxvset -g mydg set makedev=on myvset2
```

The next example sets the `compvol_access=read-write` attribute on the volume set, `myvset2`:

```
# vxvset -g mydg set compvol_access=read-write myvset2
```

The final example removes raw device node access for the volume set, `myvset2`:

```
# vxvset -g mydg set makedev=off myvset2
```


Multi-volume file systems

This chapter includes the following topics:

- [About multi-volume file systems](#)
- [About volume types](#)
- [Features implemented using multi-volume support](#)
- [Creating multi-volume file systems](#)
- [Converting a single volume file system to a multi-volume file system](#)
- [Adding a volume to and removing a volume from a multi-volume file system](#)
- [Volume encapsulation](#)
- [Reporting file extents](#)
- [Load balancing](#)
- [Converting a multi-volume file system to a single volume file system](#)

About multi-volume file systems

Veritas File System (VxFS) provides support for multi-volume file systems when used in conjunction with the Veritas Volume Manager. Using multi-volume support (MVS), a single file system can be created over multiple volumes, each volume having its own properties. For example, it is possible to place metadata on mirrored storage while placing file data on better-performing volume types such as RAID-1+0 (striped and mirrored). The volume must be in the same disk group as the volume set, and it cannot already be a member of another volume set.

The MVS feature also allows file systems to reside on different classes of devices, so that a file system can be supported from both inexpensive disks and from

expensive arrays. Using the MVS administrative interface, you can control which data goes on which volume types.

Note: Multi-volume file system support is available only on file systems using disk layout Version 7 or later.

About volume types

Veritas File System (VxFS) utilizes two types of volumes, one of which contains only data, referred to as `dataonly`, and the other of which can contain metadata or data, referred to as `metadataok`.

Data refers to direct extents, which contain user data, of regular files and named data streams in a file system.

Metadata refers to all extents that are not regular file or named data stream extents. This includes certain files that appear to be regular files, but are not, such as the File Change Log file.

A volume availability flag is set to specify if a volume is `dataonly` or `metadataok`. The volume availability flag can be set, cleared, and listed with the `fsvoladm` command.

See the `fsvoladm(1M)` manual page.

Features implemented using multi-volume support

The following features can be implemented using multi-volume support:

- Controlling where files are stored can be selected at multiple levels so that specific files or file hierarchies can be assigned to different volumes. This functionality is available in the Veritas File System `SmartTier` feature. See [“About SmartTier”](#) on page 724.
- Placing the VxFS intent log on its own volume to minimize disk head movement and thereby increase performance. This functionality can be used to migrate from the Veritas QuickLog™ feature.
- Separating Storage Checkpoints so that data allocated to a Storage Checkpoint is isolated from the rest of the file system.
- Separating metadata from file data.
- Encapsulating volumes so that a volume appears in the file system as a file. This is particularly useful for databases that are running on raw volumes.

- Guaranteeing that a dataonly volume being unavailable does not cause a metadataok volume to be unavailable.

See “[Volume availability](#)” on page 711.

To use the multi-volume file system features, Veritas Volume Manager must be installed and the volume set feature must be accessible. The volume set feature is separately licensed.

Volume availability

MVS guarantees that a dataonly volume being unavailable does not cause a metadataok volume to be unavailable. This allows you to mount a multi-volume file system even if one or more component dataonly volumes are missing.

The volumes are separated by whether metadata is allowed on the volume. An I/O error on a dataonly volume does not affect access to any other volumes. All VxFS operations that do not access the missing `dataonly` volume function normally.

Some VxFS operations that do not access the missing `dataonly` volume and function normally include the following:

- Mounting the multi-volume file system, regardless if the file system is read-only or read/write.
- Kernel operations.
- Performing a `fsck` replay. Logged writes are converted to normal writes if the corresponding volume is `dataonly`.
- Performing a full `fsck`.
- Using all other commands that do not access data on a missing volume.

Some operations that could fail if a `dataonly` volume is missing include:

- Reading or writing file data if the file's data extents were allocated from the missing `dataonly` volume.
- Using the `vxdump` command.

Volume availability is supported only on a file system with disk layout Version 7 or later.

Note: Do not mount a multi-volume system with the `ioerror=disable` or `ioerror=wdisable` mount options if the volumes have different availability properties. Symantec recommends the `ioerror=mdisable` mount option both for cluster mounts and for local mounts.

Creating multi-volume file systems

When a multi-volume file system is created, all volumes are `dataonly`, except volume zero, which is used to store the file system's metadata. The volume availability flag of volume zero cannot be set to `dataonly`.

As metadata cannot be allocated from `dataonly` volumes, enough metadata space should be allocated using `metadataok` volumes. The "file system out of space" error occurs if there is insufficient metadata space available, even if the `df` command shows that there is free space in the file system. The `fsvoladm` command can be used to see the free space in each volume and set the availability flag of the volume.

Unless otherwise specified, VxFS commands function the same on multi-volume file systems as the commands do on single-volume file systems.

Example of creating a multi-volume file system

The following procedure is an example of creating a multi-volume file system.

To create a multi-volume file system

- 1 After a volume set is created, create a VxFS file system by specifying the volume set name as an argument to `mkfs`:

```
# mkfs -F vxfs /dev/vx/rdisk/rootdg/myvset
version 9 layout
327680 sectors, 163840 blocks of size 1024,
log size 1024 blocks largefiles supported
```

After the file system is created, VxFS allocates space from the different volumes within the volume set.

- 2 List the component volumes of the volume set using of the `fsvoladm` command:

```
# mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
# fsvoladm -H list /mnt1
```

devid	size	used	avail	name
0	20 GB	10 GB	10 GB	vol1
1	30 TB	10 TB	20 TB	vol2

- 3 Add a new volume by adding the volume to the volume set, then adding the volume to the file system:

```
# vxassist -g dg1 make vol5 50m
# vxvset -g dg1 addvol myvset vol5
# fsvoladm add /mnt1 vol5 50m
# fsvoladm -H list /mnt1
```

devid	size	used	avail	name
0	20 GB	10 GB	10 GB	vol1
1	30 TB	10 TB	20 TB	vol2

- 4 List the volume availability flags using the `fsvoladm` command:

```
# fsvoladm queryflags /mnt1
```

volname	flags
vol1	metadataok
vol2	dataonly
vol3	dataonly
vol4	dataonly
vol5	dataonly

- 5 Increase the metadata space in the file system using the `fsvoladm` command:

```
# fsvoladm clearflags dataonly /mnt1 vol2
# fsvoladm queryflags /mnt1
```

volname	flags
vol1	metadataok
vol2	metadataok
vol3	dataonly
vol4	dataonly
vol5	dataonly

Converting a single volume file system to a multi-volume file system

The following procedure converts a traditional, single volume file system, `/mnt1`, on a single volume `vol1` in the diskgroup `dg1` to a multi-volume file system.

To convert a single volume file system

- 1 Determine the version of the volume's diskgroup:

```
# vxdg list dg1 | grep version: | awk '{ print $2 }'  
105
```

- 2 If the version is less than 110, upgrade the diskgroup:

```
# vxdg upgrade dg1
```

- 3 Determine the disk layout version of the file system:

```
# vxupgrade /mnt1  
Version 6
```

- 4 If the disk layout version is 6, upgrade to Version 7:

```
# vxupgrade -n 7 /mnt1
```

- 5 Unmount the file system:

```
# umount /mnt1
```

- 6 Convert the volume into a volume set:

```
# vxvset -g dg1 make vset1 vol1
```

- 7 Edit the `/etc/vfstab` file to replace the volume device name, `vol1`, with the volume set name, `vset1`.

- 8 Mount the file system:

```
# mount -F vxfs /dev/vx/dsk/dg1/vset1 /mnt1
```

- 9 As necessary, create and add volumes to the volume set:

```
# vxassist -g dg1 make vol2 256M  
# vxvset -g dg1 addvol vset1 vol2
```

- 10 Set the placement class tags on all volumes that do not have a tag:

```
# vxassist -g dg1 settag vol1 vxfs.placement_class.tier1  
# vxassist -g dg1 settag vol2 vxfs.placement_class.tier2
```

Adding a volume to and removing a volume from a multi-volume file system

Use the `fsvoladm` command to perform the following functions:

- [Adding a volume to a multi-volume file system](#)
- [Removing a volume from a multi-volume file system](#)

Use the `fsck` command to perform the following function:

- [Forcibly removing a volume in a multi-volume file system](#)

Use the `vxassist` command to perform the following function:

- [Moving volume 0 in a multi-volume file system](#)

Adding a volume to a multi-volume file system

Use the `fsvoladm add` command to add a volume to a multi-volume file system.

To add a volume to a multi-volume file system

- ◆ Add a new volume to a multi-volume file system:

```
# fsvoladm add /mnt1 vol2 256m
```

Removing a volume from a multi-volume file system

Use the `fsvoladm remove` command to remove a volume from a multi-volume file system. The `fsvoladm remove` command fails if the volume being removed is the only volume in any allocation policy.

To remove a volume from a multi-volume file system

- ◆ Remove a volume from a multi-volume file system:

```
# fsvoladm remove /mnt1 vol2
```

Forcibly removing a volume in a multi-volume file system

If you must forcibly remove a volume from a file system, such as if a volume is permanently destroyed and you want to clean up the dangling pointers to the lost volume, use the `fsck -o zapvol=volname` command. The `zapvol` option performs a full file system check and zaps all inodes that refer to the specified volume. The `fsck` command prints the inode numbers of all files that the command destroys; the file names are not printed. The `zapvol` option only affects regular files if used

on a `dataonly` volume. However, it could destroy structural files if used on a `metadataok` volume, which can make the file system unrecoverable. Therefore, the `zapvol` option should be used with caution on `metadataok` volumes.

Moving volume 0 in a multi-volume file system

Volume 0 in a multi-volume file system cannot be removed from the file system, but you can move volume 0 to different storage using the `vxassist move` command. The `vxassist` command creates any necessary temporary mirrors and cleans up the mirrors at the end of the operation.

To move volume 0

- ◆ Move volume 0:

```
# vxassist -g mydg move vol1 \!mydg
```

Volume encapsulation

Multi-volume file system support enables the ability to encapsulate an existing raw volume and make the volume contents appear as a file in the file system.

Encapsulating a volume involves the following actions:

- Adding the volume to an existing volume set.
- Adding the volume to the file system using `fsvoladm`.

Encapsulating a volume

The following example illustrates how to encapsulate a volume.

To encapsulate a volume

1 List the volumes:

```
# vxvset -g dg1 list myvset
VOLUME    INDEX    LENGTH    STATE    CONTEXT
vol1       0         102400    ACTIVE   -
vol2       1         102400    ACTIVE   -
```

The volume set has two volumes.

2 Create a third volume and copy the passwd file to the third volume:

```
# vxassist -g dg1 make dbvol 100m
# dd if=/etc/passwd of=/dev/vx/rdisk/rootdg/dbvol count=1
1+0 records in
1+0 records out
```

The third volume will be used to demonstrate how the volume can be accessed as a file, as shown later.

3 Create a file system on the volume set:

```
# mkfs -F vxfs /dev/vx/rdisk/rootdg/myvset
version 9 layout
204800 sectors, 102400 blocks of size 1024,
log size 1024 blocks
largefiles supported
```

4 Mount the volume set:

```
# mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1

# mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
```

5 Add the new volume to the volume set:

```
# vxvset -g dg1 addvol myvset dbvol
```

6 Encapsulate dbvol:

```
# fsvoladm encapsulate /mnt1/dbfile dbvol 100m
# ls -l /mnt1/dbfile
-rw----- 1 root other 104857600 May 22 11:30 /mnt1/dbfile
```

7 Examine the contents of dbfile to see that it can be accessed as a file:

```
# head -2 /mnt1/dbfile
root:x:0:1:Super-User:/:/sbin/sh
daemon:x:1:1:/:/
```

The passwd file that was written to the raw volume is now visible in the new file.

Note: If the encapsulated file is changed in any way, such as if the file is extended, truncated, or moved with an allocation policy or resized volume, or the volume is encapsulated with a bias, the file cannot be de-encapsulated.

Deencapsulating a volume

The following example illustrates how to deencapsulate a volume.

To deencapsulate a volume

1 List the volumes:

```
# vxvset -g dg1 list myvset
VOLUME    INDEX    LENGTH    STATE    CONTEXT
vol1       0         102400    ACTIVE   -
vol2       1         102400    ACTIVE   -
dbvol      2         102400    ACTIVE   -
```

The volume set has three volumes.

2 Deencapsulate dbvol:

```
# fsvoladm deencapsulate /mnt1/dbfile
```

Reporting file extents

Multi-volume file system feature provides the capability for file-to-volume mapping and volume-to-file mapping via the `fsmmap` and `fsvmap` commands. The `fsmmap`

`command` reports the volume name, logical offset, and size of data extents, or the volume name and size of indirect extents associated with a file on a multi-volume file system. The `fsvmap` command maps volumes to the files that have extents on those volumes.

See the `fsmmap(1M)` and `fsvmap(1M)` manual pages.

The `fsmmap` command requires `open()` permission for each file or directory specified. Root permission is required to report the list of files with extents on a particular volume.

Examples of reporting file extents

The following examples show typical uses of the `fsmmap` and `fsvmap` commands.

Using the `fsmmap` command

- ◆ Use the `find` command to descend directories recursively and run `fsmmap` on the list of files:

```
# find . | fsmmap -
Volume  Extent Type  File
vol2    Data             ./file1
vol1    Data             ./file2
```

Using the fsvmap command

- 1 Report the extents of files on multiple volumes:

```
# fsvmap /dev/vx/rdsk/fstest/testvset vol1 vol2
vol1 /.
vol1 /ns2
vol1 /ns3
vol1 /file1
vol2 /file1
vol2 /file2
```

- 2 Report the extents of files that have either data or metadata on a single volume in all Storage Checkpoints, and indicate if the volume has file system metadata:

```
# fsvmap -mvC /dev/vx/rdsk/fstest/testvset vol1
Meta Structural vol1 //volume has filesystem metadata//
Data UNNAMED vol1 /.
Data UNNAMED vol1 /ns2
Data UNNAMED vol1 /ns3
Data UNNAMED vol1 /file1
Meta UNNAMED vol1 /file1
```

Load balancing

An allocation policy with the balance allocation order can be defined and assigned to files that must have their allocations distributed at random between a set of specified volumes. Each extent associated with these files are limited to a maximum size that is defined as the required chunk size in the allocation policy. The distribution of the extents is mostly equal if none of the volumes are full or disabled.

Load balancing allocation policies can be assigned to individual files or for all files in the file system. Although intended for balancing data extents across volumes, a load balancing policy can be assigned as a metadata policy if desired, without any restrictions.

Note: If a file has both a fixed extent size set and an allocation policy for load balancing, certain behavior can be expected. If the chunk size in the allocation policy is greater than the fixed extent size, all extents for the file are limited by the chunk size. For example, if the chunk size is 16 MB and the fixed extent size is 3 MB, then the largest extent that satisfies both the conditions is 15 MB. If the fixed extent size is larger than the chunk size, all extents are limited to the fixed extent size. For example, if the chunk size is 2 MB and the fixed extent size is 3 MB, then all extents for the file are limited to 3 MB.

Defining and assigning a load balancing allocation policy

The following example defines a load balancing policy and assigns the policy to the file, `/mnt/file.db`.

To define and assign the policy

- 1 Define the policy by specifying the `-o balance` and `-c` options:

```
# fsapadm define -o balance -c 2m /mnt loadbal vol1 vol2 vol3 vol4
```

- 2 Assign the policy:

```
# fsapadm assign /mnt/filedb loadbal meta
```

Rebalancing extents

Extents can be rebalanced by strictly enforcing the allocation policy. Rebalancing is generally required when volumes are added or removed from the policy or when the chunk size is modified. When volumes are removed from the volume set, any extents on the volumes being removed are automatically relocated to other volumes within the policy.

The following example redefines a policy that has four volumes by adding two new volumes, removing an existing volume, and enforcing the policy for rebalancing.

To rebalance extents

- 1 Define the policy by specifying the `-o balance` and `-c` options:

```
# fsapadm define -o balance -c 2m /mnt loadbal vol1 vol2 vol4 \
vol5 vol6
```

- 2 Enforce the policy:

```
# fsapadm enforcefile -f strict /mnt/filedb
```

Converting a multi-volume file system to a single volume file system

Because data can be relocated among volumes in a multi-volume file system, you can convert a multi-volume file system to a traditional, single volume file system by moving all file system data onto a single volume. Such a conversion is useful to users who would like to try using a multi-volume file system or SmartTier, but are not committed to using a multi-volume file system permanently.

See [“About SmartTier”](#) on page 724.

There are three restrictions to this operation:

- The single volume must be the first volume in the volume set
- The first volume must have sufficient space to hold all of the data and file system metadata
- The volume cannot have any allocation policies that restrict the movement of data

The following procedure converts an existing multi-volume file system, `/mnt1`, of the volume set `vset1`, to a single volume file system, `/mnt1`, on volume `vol1` in diskgroup `dg1`.

Note: Steps 5, 6, 7, and 8 are optional, and can be performed if you prefer to remove the wrapper of the volume set object.

Converting to a single volume file system

- 1 Determine if the first volume in the volume set, which is identified as device number 0, has the capacity to receive the data from the other volumes that will be removed:

```
# df /mnt1
/mnt1  (/dev/vx/dsk/dg1/vol1):16777216 blocks  3443528 files
```

- 2 If the first volume does not have sufficient capacity, grow the volume to a sufficient size:

```
# fsvoladm resize /mnt1 vol1 150g
```

- 3 Remove all existing allocation policies:

```
# fsppadm unassign /mnt1
```

- 4 Remove all volumes except the first volume in the volume set:

```
# fsvoladm remove /mnt1 vol2
# vxvset -g dg1 rmvol vset1 vol2
# fsvoladm remove /mnt1 vol3
# vxvset -g dg1 rmvol vset1 vol3
```

Before removing a volume, the file system attempts to relocate the files on that volume. Successful relocation requires space on another volume, and no allocation policies can be enforced that pin files to that volume. The time for the command to complete is proportional to the amount of data that must be relocated.

- 5 Unmount the file system:

```
# umount /mnt1
```

- 6 Remove the volume from the volume set:

```
# vxvset -g dg1 rmvol vset1 vol1
```

- 7 Edit the `/etc/vfstab` file to replace the volume set name, `vset1`, with the volume device name, `vol1`.

- 8 Mount the file system:

```
# mount -F vxfs /dev/vx/dsk/dg1/vol1 /mnt1
```

Administering SmartTier

This chapter includes the following topics:

- [About SmartTier](#)
- [Supported SmartTier document type definitions](#)
- [Placement classes](#)
- [Administering placement policies](#)
- [File placement policy grammar](#)
- [File placement policy rules](#)
- [Calculating I/O temperature and access temperature](#)
- [Multiple criteria in file placement policy rule statements](#)
- [File placement policy rule and statement ordering](#)
- [File placement policies and extending files](#)
- [Using SmartTier with solid state disks](#)

About SmartTier

Veritas File System (VxFS) uses multi-tier online storage by way of the SmartTier feature, which functions on top of multi-volume file systems. Multi-volume file systems are file systems that occupy two or more virtual volumes. The collection of volumes is known as a volume set. A volume set is made up of disks or disk array LUNs belonging to a single Veritas Volume Manager (VxVM) disk group. A multi-volume file system presents a single name space, making the existence of multiple volumes transparent to users and applications. Each volume retains a

separate identity for administrative purposes, making it possible to control the locations to which individual files are directed.

See “[About multi-volume file systems](#)” on page 709.

Note: Some of the commands have changed or been removed between the 4.1 release and the 6.0 release to make placement policy management more user-friendly. The following commands have been removed: `fsrpadm`, `fsmove`, and `fssweep`. The output of the `queryfile`, `queryfs`, and `list` options of the `fsapadm` command now print the allocation order by name instead of number.

In the previous VxFS 5.x releases, SmartTier was known as Dynamic Storage Tiering.

SmartTier allows administrators of multi-volume VxFS file systems to manage the placement of files and the placement of portions of files on individual volumes in a volume set by defining placement policies. Placement policies control both initial file location and the circumstances under which existing files are relocated. These placement policies cause the files to which they apply to be created and extended on specific subsets of a file system's volume set, known as placement classes. The files are relocated to volumes in other placement classes when they meet the specified naming, timing, access rate, and storage capacity-related conditions.

You make a VxVM volume part of a placement class by associating a volume tag with it. For file placement purposes, VxFS treats all of the volumes in a placement class as equivalent, and balances space allocation across them. A volume may have more than one tag associated with it. If a volume has multiple tags, the volume belongs to multiple placement classes and is subject to allocation and relocation policies that relate to any of the placement classes. Multiple tagging should be used carefully.

See “[Placement classes](#)” on page 726.

VxFS imposes no capacity, performance, availability, or other constraints on placement classes. Any volume may be added to any placement class, no matter what type the volume has nor what types other volumes in the class have. However, a good practice is to place volumes of similar I/O performance and availability in the same placement class.

The *Using SmartTier* Symantec Yellow Book provides additional information regarding the SmartTier feature, including the value of SmartTier and best practices for using SmartTier. You can download *Using SmartTier* from the following Web page:

<http://www.symantec.com/enterprise/yellowbooks/index.jsp>

Supported SmartTier document type definitions

Table 33-1 describes which releases of Veritas File System (VxFS) support specific SmartTier document type definitions (DTDs).

Table 33-1 Supported SmartTier document type definitions

VxFS Version	DTD Version	
	1.0	1.1
5.0	Supported	Not supported
5.1	Supported	Supported
5.1 SP1	Supported	Supported
6.0	Supported	Supported

Placement classes

A placement class is a SmartTier attribute of a given volume in a volume set of a multi-volume file system. This attribute is a character string, and is known as a volume tag. A volume can have different tags, one of which can be the placement class. The placement class tag makes a volume distinguishable by SmartTier.

Volume tags are organized as hierarchical name spaces in which periods separate the levels of the hierarchy. By convention, the uppermost level in the volume tag hierarchy denotes the Veritas Storage Foundation Cluster File System High Availability component or application that uses a tag, and the second level denotes the tag's purpose. SmartTier recognizes volume tags of the form

`vxfs.placement_class.class_name`. The prefix `vxfs` identifies a tag as being associated with VxFS. The `placement_class` string identifies the tag as a file placement class that SmartTier uses. The `class_name` string represents the name of the file placement class to which the tagged volume belongs. For example, a volume with the tag `vxfs.placement_class.tier1` belongs to placement class `tier1`. Administrators use the `vxassist` command to associate tags with volumes.

See the `vxassist(1M)` manual page.

VxFS policy rules specify file placement in terms of placement classes rather than in terms of individual volumes. All volumes that belong to a particular placement class are interchangeable with respect to file creation and relocation operations. Specifying file placement in terms of placement classes rather than in terms of specific volumes simplifies the administration of multi-tier storage.

The administration of multi-tier storage is simplified in the following ways:

- Adding or removing volumes does not require a file placement policy change. If a volume with a tag value of `vxfs.placement_class.tier2` is added to a file system's volume set, all policies that refer to `tier2` immediately apply to the newly added volume with no administrative action. Similarly, volumes can be evacuated, that is, have data removed from them, and be removed from a file system without a policy change. The active policy continues to apply to the file system's remaining volumes.
- File placement policies are not specific to individual file systems. A file placement policy can be assigned to any file system whose volume set includes volumes tagged with the tag values (placement classes) named in the policy. This property makes it possible for data centers with large numbers of servers to define standard placement policies and apply them uniformly to all servers with a single administrative action.

Tagging volumes as placement classes

The following example tags the `vsavola` volume as placement class `tier1`, `vsavolb` as placement class `tier2`, `vsavolc` as placement class `tier3`, and `vsavold` as placement class `tier4` using the `vxassist settag` command.

To tag volumes

- ◆ Tag the volumes as placement classes:

```
# vxassist -g cfsdg settag vsavola vxfs.placement_class.tier1
# vxassist -g cfsdg settag vsavolb vxfs.placement_class.tier2
# vxassist -g cfsdg settag vsavolc vxfs.placement_class.tier3
# vxassist -g cfsdg settag vsavold vxfs.placement_class.tier4
```

Listing placement classes

Placement classes are listed using the `vxassist listtag` command.

See the `vxassist(1M)` manual page.

The following example lists all volume tags, including placement classes, set on a volume `vsavola` in the diskgroup `cfsdg`.

To list placement classes

- ◆ List the volume tags, including placement classes:

```
# vxassist -g cfsdg listtag vsavola
```

Administering placement policies

A VxFS file placement policy document contains rules by which VxFS creates, relocates, and deletes files, but the placement policy does not refer to specific file systems or volumes. You can create a file system's active file placement policy by assigning a placement policy document to the file system via the `fsppadm` command or the GUI.

See the `fsppadm(1M)` manual page.

At most, one file placement policy can be assigned to a VxFS file system at any time. A file system may have no file placement policy assigned to it, in which case VxFS allocates space for new files according to its own internal algorithms.

In systems with Storage Foundation Management Server (SFMS) software installed, file placement policy information is stored in the SFMS database. The SFMS database contains both XML policy documents and lists of hosts and file systems for which each document is the current active policy. When a policy document is updated, SFMS can assign the updated document to all file systems whose current active policies are based on that document. By default, SFMS does not update file system active policies that have been created or modified locally, that is by the hosts that control the placement policies' file systems. If a SFMS administrator forces assignment of a placement policy to a file system, the file system's active placement policy is overwritten and any local changes that had been made to the placement policy are lost.

You can view sample placement policies in the `/opt/VRTSvxfs/etc` directory. These sample placement policies are installed as part of the VxFS package installation.

Assigning a placement policy

The following example uses the `fsppadm assign` command to assign the file placement policy represented in the XML policy document `/tmp/policy1.xml` for the file system at mount point `/mnt1`.

To assign a placement policy

- ◆ Assign a placement policy to a file system:

```
# fsppadm assign /mnt1 /tmp/policy1.xml
```

Unassigning a placement policy

The following example uses the `fsppadm unassign` command to unassign the active file placement policy from the file system at mount point `/mnt1`.

To unassign a placement policy

- ◆ Unassign the placement policy from a file system:

```
# fsppadm unassign /mnt1
```

Analyzing the space impact of enforcing a placement policy

The following example uses the `fsppadm analyze` command to analyze the impact if the enforce operation is performed on the file placement policy represented in the XML policy document `/tmp/policy1.xml` for the mount point `/mnt1`. The command builds the I/O temperature database if necessary.

To analyze the space impact of enforcing a placement policy

- ◆ Analyze the impact of enforcing the file placement policy represented in the XML policy document `/tmp/policy1.xml` for the mount point `/mnt1`:

```
# fsppadm analyze -F /tmp/policy1.xml -i /mnt1
```

Querying which files will be affected by enforcing a placement policy

The following example uses the `fsppadm query` command to generate a list of files that will be affected by enforcing a placement policy. The command provides details about where the files currently reside, to where the files will be relocated, and which rule in the placement policy applies to the files.

To query which files will be affected by enforcing a placement policy

- ◆ Query the files that will be affected:

```
# fsppadm query /mnt1/dir1/dir2 /mnt2 /mnt1/dir3
```

Enforcing a placement policy

Enforcing a placement policy for a file system requires that the policy be assigned to the file system. You must assign a placement policy before it can be enforced.

See [“Assigning a placement policy”](#) on page 728.

Enforce operations are logged in a hidden file, `.__fsppadm_enforce.log`, in the `lost+found` directory of the mount point. This log file contains details such as files' previous locations, the files' new locations, and the reasons for the files' relocations. The enforce operation creates the `.__fsppadm_enforce.log` file if the file does not exist. The enforce operation appends the file if the file already exists. The `.__fsppadm_enforce.log` file can be backed up or removed as with a normal file.

You can specify the `-F` option to specify a placement policy other than the existing active placement policy. This option can be used to enforce the rules given in the specified placement policy for maintenance purposes, such as for reclaiming a LUN from the file system.

You can specify the `-p` option to specify the number of concurrent threads to be used to perform the `fsppadm` operation. You specify the `io_nice` parameter as an integer between 1 and 100, with 50 being the default value. A value of 1 specifies 1 slave and 1 master thread per mount. A value of 50 specifies 16 slaves and 1 master thread per mount. A value of 100 specifies 32 slaves and 1 master thread per mount.

You can specify the `-C` option so that the `fsppadm` command processes only those files that have some activity stats logged in the File Change Log (FCL) file during the period specified in the placement policy. You can use the `-C` option only if the policy's `ACCESSTEMP` or `IOTEMP` elements use the `Prefer` criteria.

You can specify the `-T` option to specify the placement classes that contain files for the `fsppadm` command to sweep and relocate selectively. You can specify the `-T` option only if the policy uses the `Prefer` criteria for `IOTEMP`.

See the `fsppadm(1M)` manual page.

The following example uses the `fsppadm enforce` command to enforce the file placement policy for the file system at mount point `/mnt1`, and includes the access time, modification time, and file size of the specified paths in the report, `/tmp/report`.

To enforce a placement policy

- ◆ Enforce a placement policy for a file system:

```
# fsppadm enforce -a -r /tmp/report /mnt1
Current Current Relocated Relocated
Class Volume Class Volume Rule File
tier3 vole tier3 vole a_to_z /mnt1/mds1/d1/file1
tier3 vole tier3 vole a_to_z /mnt1/mds1/d1/file2
tier3 vole tier3 vole a_to_z /mnt1/mds1/d1/d2/file3
tier3 volf tier3 volf a_to_z /mnt1/mds1/d1/d2/file4
.
.
.
Sweep path : /mnt1
Files moved : 42
KB moved : 1267
```

Tier Name	Size (KB)	Free Before (KB)	Free After (KB)
tier4	524288	524256	524256
tier3	524288	522968	522968
tier2	524288	524256	524256
tier1	524288	502188	501227

Validating a placement policy

The following example uses the `fsppadm validate` command to validate the placement policy `policy.xml` against all mounted file systems.

To validate a placement policy against all mounted file systems

- ◆ Validate the placement policy:

```
# fsppadm validate /tmp/policy.xml
```

File placement policy grammar

VxFS allocates and relocates files within a multi-volume file system based on properties in the file system metadata that pertains to the files. Placement decisions may be based on file name, directory of residence, time of last access, access frequency, file size, and ownership. An individual file system's criteria for allocating and relocating files are expressed in the file system's file placement policy.

A VxFS file placement policy defines the desired placement of sets of files on the volumes of a VxFS multi-volume file system. A file placement policy specifies the placement classes of volumes on which files should be created, and where and under what conditions the files should be relocated to volumes in alternate placement classes or deleted. You can create file placement policy documents, which are XML text files, using an XML editor, a text editor, or Veritas Operations Manager (VOM).

See the `/opt/VRTSvxfs/etc/placement_policy.dtd` file for the overall structure of a placement policy.

File placement policy rules

A VxFS file placement policy consists of one or more rules. Each rule applies to one or more files. The files to which a rule applies are designated in one or more `SELECT` statements. A `SELECT` statement designates files according to one or more of four properties: their names or naming patterns, the directories in which they reside, their owners' user names, and their owners' group names.

A file may be designated by more than one rule. For example, if one rule designates files in directory `/dir`, and another designates files owned by `user1`, a file in `/dir` that is owned by `user1` is designated by both rules. Only the rule that appears first in the placement policy applies to the file; subsequent rules are ignored.

You can define placement policies that do not encompass the entire file system name space. When a file that is not designated by any rule in its file system's active placement policy is created, VxFS places the file according to its own internal algorithms. To maintain full control over file placement, include a catchall rule at the end of each placement policy document with a `SELECT` statement that designates files by the naming pattern `*`. Such a rule designates all files that have not been designated by the rules appearing earlier in the placement policy document.

Two types of rules exist: `data` and `ckpt`. The `data` rule type allows SmartTier to relocate normal data files. The `ckpt` rule type allows SmartTier to relocate Storage Checkpoints. You specify the rule type by setting the `Flags` attribute for the rule.

SELECT statement

The VxFS placement policy rule `SELECT` statement designates the collection of files to which a rule applies.

The following XML snippet illustrates the general form of the `SELECT` statement:

```
<SELECT>
  <DIRECTORY Flags="directory_flag_value"> value
</DIRECTORY>
```



```

<PATTERN Flags="pattern_flag_value"> value </PATTERN>
<USER> value </USER>
<GROUP> value </GROUP>
</SELECT>

```

A **SELECT** statement may designate files by using the following selection criteria:

<DIRECTORY> A full path name relative to the file system mount point. The `Flags="directory_flag_value"` XML attribute must have a value of `nonrecursive`, denoting that only files in the specified directory are designated, or a value of `recursive`, denoting that files in all subdirectories of the specified directory are designated. The `Flags` attribute is mandatory.

The **<DIRECTORY>** criterion is optional, and may be specified more than once.

<PATTERN> Either an exact file name or a pattern using a single wildcard character (*). For example, the pattern "abc*" denotes all files whose names begin with "abc". The pattern "abc.*" denotes all files whose names are exactly "abc" followed by a period and any extension. The pattern "*abc" denotes all files whose names end in "abc", even if the name is all or part of an extension. The pattern "*.abc" denotes files of any name whose name extension (following the period) is "abc". The pattern "ab*c" denotes all files whose names start with "ab" and end with "c". The first "*" character is treated as a wildcard, while any subsequent "*" characters are treated as literal text. The pattern cannot contain "/".

The wildcard character matches any character, including ".", "?", and "[", unlike using the wildcard in a shell.

The `Flags="pattern_flag_value"` XML attribute is optional, and if specified can only have a value of `recursive`. Specify `Flags="recursive"` only if the pattern is a directory. If `Flags` is not specified, the default attribute value is `nonrecursive`. If `Flags="recursive"` is specified, the enclosing selection criteria selects all files in any component directory that is anywhere below the directory specified by **<DIRECTORY>** if the component directory matches the pattern and either of the following is true:

- **<DIRECTORY>** is specified and has the recursive flag.
- **<DIRECTORY>** is not specified and the directory is anywhere in the file system.

If the pattern contains the wildcard character (*), wildcard character matching is performed.

The **<PATTERN>** criterion is optional, and may be specified more than once. Only one value can be specified per **<PATTERN>** element.

<USER>	User name of the file's owner. The user number cannot be specified in place of the name. The <USER> criterion is optional, and may be specified more than once.
<GROUP>	Group name of the file's owner. The group number cannot be specified in place of the group name. The <GROUP> criterion is optional, and may be specified more than once.

One or more instances of any or all of the file selection criteria may be specified within a single `SELECT` statement. If two or more selection criteria of different types are specified in a single statement, a file must satisfy one criterion of each type to be selected.

In the following example, only files that reside in either the `ora/db` or the `crash/dump` directory, and whose owner is either `user1` or `user2` are selected for possible action:

```
<SELECT>
  <DIRECTORY Flags="nonrecursive">ora/db</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">crash/dump</DIRECTORY>
  <USER>user1</USER>
  <USER>user2</USER>
</SELECT>
```

A rule may include multiple `SELECT` statements. If a file satisfies the selection criteria of one of the `SELECT` statements, it is eligible for action.

In the following example, any files owned by either `user1` or `user2`, no matter in which directories they reside, as well as all files in the `ora/db` or `crash/dump` directories, no matter which users own them, are eligible for action:

```
<SELECT>
  <DIRECTORY Flags="nonrecursive">ora/db</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">crash/dump</DIRECTORY>
</SELECT>
<SELECT>
  <USER>user1</USER>
  <USER>user2</USER>
</SELECT>
```

When VxFS creates new files, VxFS applies active placement policy rules in the order of appearance in the active placement policy's XML source file. The first rule in which a `SELECT` statement designates the file to be created determines the file's placement; no later rules apply. Similarly, VxFS scans the active policy rules on behalf of each file when relocating files, stopping the rules scan when it reaches

the first rule containing a `SELECT` statement that designates the file. This behavior holds true even if the applicable rule results in no action. Take for example a policy rule that indicates that `.dat` files inactive for 30 days should be relocated, and a later rule indicates that `.dat` files larger than 10 megabytes should be relocated. A 20 megabyte `.dat` file that has been inactive for 10 days will not be relocated because the earlier rule applied. The later rule is never scanned.

A placement policy rule's action statements apply to all files designated by any of the rule's `SELECT` statements. If an existing file is not designated by a `SELECT` statement in any rule of a file system's active placement policy, then SmartTier does not relocate or delete the file. If an application creates a file that is not designated by a `SELECT` statement in a rule of the file system's active policy, then VxFS places the file according to its own internal algorithms. If this behavior is inappropriate, the last rule in the policy document on which the file system's active placement policy is based should specify `<PATTERN>*</PATTERN>` as the only selection criterion in its `SELECT` statement, and a `CREATE` statement naming the desired placement class for files not selected by other rules.

CREATE statement

A `CREATE` statement in a file placement policy rule specifies one or more placement classes of volumes on which VxFS should allocate space for new files to which the rule applies at the time the files are created. You can specify only placement classes, not individual volume names, in a `CREATE` statement.

A file placement policy rule may contain at most one `CREATE` statement. If a rule does not contain a `CREATE` statement, VxFS places files designated by the rule's `SELECT` statements according to its internal algorithms. However, rules without `CREATE` statements can be used to relocate or delete existing files that the rules' `SELECT` statements designate.

The following XML snippet illustrates the general form of the `CREATE` statement:

```
<CREATE>
  <ON Flags="flag_value">
    <DESTINATION>
      <CLASS> placement_class_name </CLASS>
      <BALANCE_SIZE Units="units_specifier"> chunk_size
    </BALANCE_SIZE>
    </DESTINATION>
    <DESTINATION> additional_placement_class_specifications
  </DESTINATION>
  </ON>
</CREATE>
```

A `CREATE` statement includes a single `<ON>` clause, in which one or more `<DESTINATION>` XML elements specify placement classes for initial file allocation in order of decreasing preference. VxFS allocates space for new files to which a rule applies on a volume in the first class specified, if available space permits. If space cannot be allocated on any volume in the first class, VxFS allocates space on a volume in the second class specified if available space permits, and so forth.

If space cannot be allocated on any volume in any of the placement classes specified, file creation fails with an `ENOSPC` error, even if adequate space is available elsewhere in the file system's volume set. This situation can be circumvented by specifying a `Flags` attribute with a value of "any" in the `<ON>` clause. If `<ON Flags="any">` is specified in a `CREATE` statement, VxFS first attempts to allocate space for new files to which the rule applies on the specified placement classes. Failing that, VxFS resorts to its internal space allocation algorithms, so file allocation does not fail unless there is no available space any-where in the file system's volume set.

The `Flags="any"` attribute differs from the catchall rule in that this attribute applies only to files designated by the `SELECT` statement in the rule, which may be less inclusive than the `<PATTERN>*</PATTERN>` file selection specification of the catchall rule.

In addition to the placement class name specified in the `<CLASS>` sub-element, a `<DESTINATION>` XML element may contain a `<BALANCE_SIZE>` sub-element. Presence of a `<BALANCE_SIZE>` element indicates that space allocation should be distributed across the volumes of the placement class in chunks of the indicated size. For example, if a balance size of one megabyte is specified for a placement class containing three volumes, VxFS allocates the first megabyte of space for a new or extending file on the first (lowest indexed) volume in the class, the second megabyte on the second volume, the third megabyte on the third volume, the fourth megabyte on the first volume, and so forth. Using the `Units` attribute in the `<BALANCE_SIZE>` XML tag, the balance size value may be specified in the following units:

bytes	Bytes
KB	Kilobytes
MB	Megabytes
GB	Gigabytes

The `<BALANCE_SIZE>` element distributes the allocation of database files across the volumes in a placement class. In principle, distributing the data in each file across multiple volumes distributes the I/O load across the volumes as well.

The `CREATE` statement in the following example specifies that files to which the rule applies should be created on the `tier1` volume if space is available, and on one of the `tier2` volumes if not. If space allocation on `tier1` and `tier2` volumes is not possible, file creation fails, even if space is available on `tier3` volumes.

```
<CREATE>
  <ON>
    <DESTINATION>
      <CLASS>tier1</CLASS>
    </DESTINATION>
    <DESTINATION>
      <CLASS>tier2</CLASS>
      <BALANCE_SIZE Units="MB">1</BALANCE_SIZE>
    </DESTINATION>
  </ON>
</CREATE>
```

The `<BALANCE_SIZE>` element with a value of one megabyte is specified for allocations on `tier2` volumes. For files allocated on `tier2` volumes, the first megabyte would be allocated on the first volume, the second on the second volume, and so forth.

RELOCATE statement

The `RELOCATE` action statement of file placement policy rules specifies an action that VxFS takes on designated files during periodic scans of the file system, and the circumstances under which the actions should be taken. The `fsppadm enforce` command is used to scan all or part of a file system for files that should be relocated based on rules in the active placement policy at the time of the scan.

See the `fsppadm(1M)` manual page.

The `fsppadm enforce` command scans file systems in path name order. For each file, VxFS identifies the first applicable rule in the active placement policy, as determined by the rules' `SELECT` statements. If the file resides on a volume specified in the `<FROM>` clause of one of the rule's `RELOCATE` statements, and if the file meets the criteria for relocation specified in the statement's `<WHEN>` clause, the file is scheduled for relocation to a volume in the first placement class listed in the `<TO>` clause that has space available for the file. The scan that results from issuing the `fsppadm enforce` command runs to completion before any files are relocated.

The following XML snippet illustrates the general form of the `RELOCATE` statement:

```
<RELOCATE>
  <FROM>
```

```

<SOURCE>
  <CLASS> placement_class_name </CLASS>
</SOURCE>
<SOURCE> additional_placement_class_specifications
</SOURCE>
</FROM>
<TO>
  <DESTINATION>
    <CLASS> placement_class_name </CLASS>
    <BALANCE_SIZE Units="units_specifier">
      chunk_size
    </BALANCE_SIZE>
  </DESTINATION>
  <DESTINATION>
    additional_placement_class_specifications
  </DESTINATION>
</TO>
<WHEN> relocation_conditions </WHEN>
</RELOCATE>

```

A **RELOCATE** statement contains the following clauses:

<FROM>	An optional clause that contains a list of placement classes from whose volumes designated files should be relocated if the files meet the conditions specified in the <WHEN> clause. No priority is associated with the ordering of placement classes listed in a <FROM> clause. If a file to which the rule applies is located on a volume in any specified placement class, the file is considered for relocation. If a RELOCATE statement contains a <FROM> clause, VxFS only considers files that reside on volumes in placement classes specified in the clause for relocation. If no <FROM> clause is present, qualifying files are relocated regardless of where the files reside.

<TO>

Indicates the placement classes to which qualifying files should be relocated. Unlike the source placement class list in a FROM clause, placement classes in a <TO> clause are specified in priority order. Files are relocated to volumes in the first specified placement class if possible, to the second if not, and so forth.

The <TO> clause of the RELOCATE statement contains a list of <DESTINATION> XML elements specifying placement classes to whose volumes VxFS relocates qualifying files. Placement classes are specified in priority order. VxFS relocates qualifying files to volumes in the first placement class specified as long as space is available. A <DESTINATION> element may contain an optional <BALANCE_SIZE> modifier sub-element. The <BALANCE_SIZE> modifier indicates that relocated files should be distributed across the volumes of the destination placement class in chunks of the indicated size. For example, if a balance size of one megabyte is specified for a placement class containing three volumes, VxFS relocates the first megabyte the file to the first (lowest indexed) volume in the class, the second megabyte to the second volume, the third megabyte to the third volume, the fourth megabyte to the first volume, and so forth. Using the Units attribute in the <BALANCE_SIZE> XML tag, the chunk value may be specified in bytes (Units="bytes"), kilobytes (Units="KB"), megabytes (Units="MB"), or gigabytes (Units="GB").

The <BALANCE_SIZE> element distributes the allocation of database files across the volumes in a placement class. In principle, distributing the data in each file across multiple volumes distributes the I/O load across the volumes as well.

<WHEN>

An optional clause that indicates the conditions under which files to which the rule applies should be relocated. Files that have been unaccessed or unmodified for a specified period, reached a certain size, or reached a specific I/O temperature or access temperature level may be relocated. If a RELOCATE statement does not contain a <WHEN> clause, files to which the rule applies are relocated unconditionally.

A <WHEN> clause may be included in a RELOCATE statement to specify that files should be relocated only if any or all of four types of criteria are met. Files can be specified for relocation if they satisfy one or more criteria.

The following are the criteria that can be specified for the <WHEN> clause:

<ACCAGE>

This criterion is met when files are inactive for a designated period or during a designated period relative to the time at which the `fsppadm enforce` command was issued.

<MODAGE>	This criterion is met when files are unmodified for a designated period or during a designated period relative to the time at which the <code>fsppadm enforce</code> command was issued.
<SIZE>	This criterion is met when files exceed or drop below a designated size or fall within a designated size range.
<IOTEMP>	This criterion is met when files exceed or drop below a designated I/O temperature, or fall within a designated I/O temperature range. A file's I/O temperature is a measure of the I/O activity against it during the period designated by the <PERIOD> element prior to the time at which the <code>fsppadm enforce</code> command was issued. See "Calculating I/O temperature and access temperature" on page 753.
<ACCESSTEMP>	This criterion is met when files exceed or drop below a specified average access temperature, or fall within a specified access temperature range. A file's access temperature is similar to its I/O temperature, except that access temperature is computed using the number of I/O requests to the file, rather than the number of bytes transferred.

Note: The use of <IOTEMP> and <ACCESSTEMP> for data placement on VxFS servers that are used as NFS servers may not be very effective due to NFS caching. NFS client side caching and the way that NFS works can result in I/O initiated from an NFS client not producing NFS server side I/O. As such, any temperature measurements in place on the server side will not correctly reflect the I/O behavior that is specified by the placement policy.

If the server is solely used as an NFS server, this problem can potentially be mitigated by suitably adjusting or lowering the temperature thresholds. However, adjusting the thresholds may not always create the desired effect. In addition, if the same mount point is used both as an NFS export as well as a local mount, the temperature-based placement decisions will not be very effective due to the NFS cache skew.

The following XML snippet illustrates the general form of the <WHEN> clause in a `RELOCATE` statement:

```
<WHEN>
  <ACCAGE Units="units_value">
    <MIN Flags="comparison_operator">
      min_access_age</MIN>
    <MAX Flags="comparison_operator">
```



```

        max_access_age</MAX>
    </ACCAGE>
    <MODAGE Units="units_value">
        <MIN Flags="comparison_operator">
            min_modification_age</MIN>
        <MAX Flags="comparison_operator">
            max_modification_age</MAX>
    </MODAGE>
    <SIZE " Units="units_value">
        <MIN Flags="comparison_operator">
            min_size</MIN>
        <MAX Flags="comparison_operator">
            max_size</MAX>
    </SIZE>
    <IOTEMP Type="read_write_preference" Prefer="temperature_preference">
        <MIN Flags="comparison_operator">
            min_I/O_temperature</MIN>
        <MAX Flags="comparison_operator">
            max_I/O_temperature</MAX>
        <PERIOD Units="days_or_hours"> days_or_hours_of_interest </PERIOD>
    </IOTEMP>
    <ACCESSTEMP Type="read_write_preference"
        Prefer="temperature_preference">
        <MIN Flags="comparison_operator">
            min_access_temperature</MIN>
        <MAX Flags="comparison_operator">
            max_access_temperature</MAX>
        <PERIOD Units="days_or_hours"> days_or_hours_of_interest </PERIOD>
    </ACCESSTEMP>
</WHEN>

```

The access age (<ACCAGE>) element refers to the amount of time since a file was last accessed. VxFS computes access age by subtracting a file's time of last access, `atime`, from the time when the `fsppadm enforce` command was issued. The <MIN> and <MAX> XML elements in an <ACCAGE> clause, denote the minimum and maximum access age thresholds for relocation, respectively. These elements are optional, but at least one must be included. Using the `Units` XML attribute, the <MIN> and <MAX> elements may be specified in the following units:

hours	Hours
days	Days. A day is considered to be 24 hours prior to the time that the <code>fsppadm enforce</code> command was issued.

Both the `<MIN>` and `<MAX>` elements require `Flags` attributes to direct their operation.

For `<MIN>`, the following `Flags` attributes values may be specified:

<code>gt</code>	The time of last access must be greater than the specified interval.
<code>eq</code>	The time of last access must be equal to the specified interval.
<code>gteq</code>	The time of last access must be greater than or equal to the specified interval.

For `<MAX>`, the following `Flags` attributes values may be specified.

<code>lt</code>	The time of last access must be less than the specified interval.
<code>lteq</code>	The time of last access must be less than or equal to the specified interval.

Including a `<MIN>` element in a `<WHEN>` clause causes VxFS to relocate files to which the rule applies that have been inactive for longer than the specified interval. Such a rule would typically be used to relocate inactive files to less expensive storage tiers. Conversely, including `<MAX>` causes files accessed within the specified interval to be relocated. It would typically be used to move inactive files against which activity had recommenced to higher performance or more reliable storage. Including both `<MIN>` and `<MAX>` causes VxFS to relocate files whose access age lies between the two.

The modification age relocation criterion, `<MODAGE>`, is similar to access age, except that files' POSIX `mtime` values are used in computations. You would typically specify the `<MODAGE>` criterion to cause relocation of recently modified files to higher performance or more reliable storage tiers in anticipation that the files would be accessed recurrently in the near future.

The file size relocation criterion, `<SIZE>`, causes files to be relocated if the files are larger or smaller than the values specified in the `<MIN>` and `<MAX>` relocation criteria, respectively, at the time that the `fsppadm enforce` command was issued. Specifying both criteria causes VxFS to schedule relocation for files whose sizes lie between the two. Using the `Units` attribute, threshold file sizes may be specified in the following units:

<code>bytes</code>	Bytes
<code>KB</code>	Kilobytes
<code>MB</code>	Megabytes

GB

Gigabytes

Specifying the I/O temperature relocation criterion

The I/O temperature relocation criterion, `<IOTEMP>`, causes files to be relocated if their I/O temperatures rise above or drop below specified values over a specified period immediately prior to the time at which the `fsppadm enforce` command was issued. A file's I/O temperature is a measure of the read, write, or total I/O activity against it normalized to the file's size. Higher I/O temperatures indicate higher levels of application activity; lower temperatures indicate lower levels. VxFS computes a file's I/O temperature by dividing the number of bytes transferred to or from it (read, written, or both) during the specified period by its size at the time that the `fsppadm enforce` command was issued.

See [“Calculating I/O temperature and access temperature”](#) on page 753.

As with the other file relocation criteria, `<IOTEMP>` may be specified with a lower threshold by using the `<MIN>` element, an upper threshold by using the `<MAX>` element, or as a range by using both. However, I/O temperature is dimensionless and therefore has no specification for units.

VxFS computes files' I/O temperatures over the period between the time when the `fsppadm enforce` command was issued and the number of days or hours in the past specified in the `<PERIOD>` element, where a day is a 24 hour period. The default unit of time is days. You can specify hours as the time unit by setting the `Units` attribute of the `<PERIOD>` element to `hours`. Symantec recommends that you specify hours only if you are using solid state disks (SSDs).

See [“Frequent SmartTier scans with solid state disks”](#) on page 765.

For example, if you issued the `fsppadm enforce` command at 2 PM on Wednesday and you want VxFS to look at file I/O activity for the period between 2 PM on Monday and 2 PM on Wednesday, which is a period of 2 days, you would specify the following `<PERIOD>` element:

```
<PERIOD> 2 </PERIOD>
```

If you instead want VxFS to look at file I/O activity between 3 hours prior to running the `fsppadm enforce` command and the time that you ran the command, you specify the following `<PERIOD>` element:

```
<PERIOD Units="hours"> 3 </PERIOD>
```

The amount of time specified in the `<PERIOD>` element should not exceed one or two weeks due to the disk space used by the File Change Log (FCL) file.

See [“About the File Change Log file”](#) on page 931.

I/O temperature is a softer measure of I/O activity than access age. With access age, a single access to a file resets the file's atime to the current time. In contrast, a file's I/O temperature decreases gradually as time passes without the file being accessed, and increases gradually as the file is accessed periodically. For example, if a new 10 megabyte file is read completely five times on Monday and `fsppadm enforce` runs at midnight, the file's two-day I/O temperature will be five and its access age in days will be zero. If the file is read once on Tuesday, the file's access age in days at midnight will be zero, and its two-day I/O temperature will have dropped to three. If the file is read once on Wednesday, the file's access age at midnight will still be zero, but its two-day I/O temperature will have dropped to one, as the influence of Monday's I/O will have disappeared.

If the intention of a file placement policy is to keep files in place, such as on top-tier storage devices, as long as the files are being accessed at all, then access age is the more appropriate relocation criterion. However, if the intention is to relocate files as the I/O load on them decreases, then I/O temperature is more appropriate.

The case for upward relocation is similar. If files that have been relocated to lower-tier storage devices due to infrequent access experience renewed application activity, then it may be appropriate to relocate those files to top-tier devices. A policy rule that uses access age with a low `<MAX>` value, that is, the interval between `fsppadm enforce` runs, as a relocation criterion will cause files to be relocated that have been accessed even once during the interval. Conversely, a policy that uses I/O temperature with a `<MIN>` value will only relocate files that have experienced a sustained level of activity over the period of interest.

Prefer attribute

You can specify a value for the `Prefer` attribute for the `<IOTEMP>` and `<ACCESSTEMP>` criteria, which gives preference to relocating files. The `Prefer` attribute can take two values: `low` or `high`. If you specify `low`, Veritas File System (VxFS) relocates the files with the lower I/O temperature before relocating the files with the higher I/O temperature. If you specify `high`, VxFS relocates the files with the higher I/O temperature before relocating the files with the lower I/O temperature. Symantec recommends that you specify a `Prefer` attribute value only if you are using solid state disks (SSDs).

See [“Prefer mechanism with solid state disks”](#) on page 764.

Different `<PERIOD>` elements may be used in the `<IOTEMP>` and `<ACCESSTEMP>` criteria of different `RELOCATE` statements within the same policy.

The following placement policy snippet gives an example of the `Prefer` criteria:

```
<RELOCATE>
...
```

```

<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high">
    <MIN Flags="gteq"> 3.4 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</RELOCATE>

```

If there are a number of files whose I/O temperature is greater than the given minimum value, the files with the higher temperature are first subject to the `RELOCATE` operation before the files with the lower temperature.

Average I/O activity criteria

The `Average` criteria allows you to specify the value of the I/O temperature as a ratio of per-file activity that occurs over the time specified by the `<PERIOD>` element compared to the overall file system activity that occurs over a longer period of time. The `<PERIOD>` element in the `RELOCATE` criteria specifies the a number of hours or days immediately before the time of the scan. During that time, the I/O statistics that are collected are used to process the files that are being scanned. Since I/O activity can change over time, collect the average I/O activity over a longer duration than the `<PERIOD>` value itself, which is by default 24 hours. Doing so lets you compute an average temperature of the whole file system. Symantec recommends that you specify an `Average` attribute value only if you are using solid state disks (SSDs).

See [“Average I/O activity with solid state disks”](#) on page 764.

The following placement policy snippet gives an example of the `Average` criteria:

```

<RELOCATE>
  ...
  <WHEN>
    <IOTEMP Type="nrbytes" Prefer="high" Average="*">
      <MIN Flags="gteq"> 1.5 </MIN>
      <PERIOD Units="hours"> 6 </PERIOD>
    </IOTEMP>
  </WHEN>
</RELOCATE>

```

In the snippet, VxFS relocates any file whose read IOTEMP over the last 6 hours is 1.5 times that of all the active files in the whole file system over the last 24 hours. This `Average` criteria is more intuitive and easier to specify than the absolute values.

The following formula computes the read IOTEMP of a given file:

$$\text{IOTEMP} = (\text{bytes of the file that are read in the PERIOD}) /$$
$$(\text{PERIOD in hours} * \text{size of the file in bytes})$$

The write and read/write IOTEMP are also computed accordingly.

The following formula computes the average read IOTEMP:

$$\text{Average IOTEMP} = (\text{bytes read of all active files in the last } h \text{ hours}) /$$
$$(h * \text{size of all the active files in bytes})$$

h is 24 hours by default. The average write and read/write IOTEMP are also computed accordingly.

In the example snippet, the value 1.5 is the multiple of average read IOTEMP over the last 24 hours across the whole file system, or rather across all of the active inodes whose activity is still available in the File Change Log (FCL) file at the time of the scan. Thus, the files' read IOTEMP activity over the last 6 hours is compared against 1.5 times that of the last 24 hours average activity to make the relocation decision. Using this method eliminates the need to give a specific number for the <IOTEMP> or <ACCESSTEMP> criteria, and instead lets you specify a multiple of the Average temperature. Keeping this averaging period longer than the specified <PERIOD> value normalizes the effects of any spikes and lulls in the file activity.

You can also use the *Average* criteria with the <ACCESSTEMP> criteria. The purpose and usage are the same.

You determine the type of the average by whether you specify the *Average* criteria with the <IOTEMP> or with the <ACCESSTEMP> criteria. The *Average* criteria can be any of the following types, depending on the criteria used:

- read Average IOTEMP
- write Average IOTEMP
- rw Average IOTEMP
- read Average ACCESSTEMP
- write Average ACCESSTEMP
- rw Average ACCESSTEMP

The default *Average* is a 24 hour average temperature, which is the total of all of the temperatures available up to the last 24 hours in the FCL file, divided by the number of files for which such I/O statistics still exist in the FCL file. You can override the number of hours by specifying the *AveragePeriod* attribute in the <PLACEMENT_POLICY> element. Symantec recommends that you specify an *AveragePeriod* attribute value only if you are using solid state disks (SSDs).

The following example statement causes the average file system activity be collected and computed over a period of 30 hours instead of the default 24 hours:

```
<PLACEMENT_POLICY Name="Policy1" Version="5.1" AveragePeriod="30">
```

RELOCATE statement examples

The following example illustrates an unconditional relocation statement, which is the simplest form of the `RELOCATE` policy rule statement:

```
<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS>tier1</CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
  </TO>
</RELOCATE>
```

The files designated by the rule's `SELECT` statement that reside on volumes in placement class `tier1` at the time the `fsppadm enforce` command executes would be unconditionally relocated to volumes in placement class `tier2` as long as space permitted. This type of rule might be used, for example, with applications that create and access new files but seldom access existing files once they have been processed. A `CREATE` statement would specify creation on `tier1` volumes, which are presumably high performance or high availability, or both. Each instantiation of `fsppadm enforce` would relocate files created since the last run to `tier2` volumes.

The following example illustrates a more comprehensive form of the `RELOCATE` statement that uses access age as the criterion for relocating files from `tier1` volumes to `tier2` volumes. This rule is designed to maintain free space on `tier1` volumes by relocating inactive files to `tier2` volumes:

```
<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS>tier1</CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
```

```

        <CLASS>tier2</CLASS>
    </DESTINATION>
</TO>
<WHEN>
    <SIZE Units="MB">
        <MIN Flags="gt">1</MIN>
        <MAX Flags="lt">1000</MAX>
    </SIZE>
    <ACCAGE Units="days">
        <MIN Flags="gt">30</MIN>
    </ACCAGE>
</WHEN>
</RELOCATE>

```

Files designated by the rule's `SELECT` statement are relocated from `tier1` volumes to `tier2` volumes if they are between 1 MB and 1000 MB in size and have not been accessed for 30 days. VxFS relocates qualifying files in the order in which it encounters them as it scans the file system's directory tree. VxFS stops scheduling qualifying files for relocation when it calculates that already-scheduled relocations would result in `tier2` volumes being fully occupied.

The following example illustrates a possible companion rule that relocates files from `tier2` volumes to `tier1` ones based on their I/O temperatures. This rule might be used to return files that had been relocated to `tier2` volumes due to inactivity to `tier1` volumes when application activity against them increases. Using I/O temperature rather than access age as the relocation criterion reduces the chance of relocating files that are not actually being used frequently by applications. This rule does not cause files to be relocated unless there is sustained activity against them over the most recent two-day period.

```

<RELOCATE>
    <FROM>
        <SOURCE>
            <CLASS>tier2</CLASS>
        </SOURCE>
    </FROM>
    <TO>
        <DESTINATION>
            <CLASS>tier1</CLASS>
        </DESTINATION>
    </TO>
    <WHEN>
        <IOTEMP Type="nrbytes">
            <MIN Flags="gt">5</MIN>

```



```

        <PERIOD>2</PERIOD>
    </IOTEMP>
</WHEN>
</RELOCATE>

```

This rule relocates files that reside on `tier2` volumes to `tier1` volumes if their I/O temperatures are above 5 for the two day period immediately preceding the issuing of the `fsppadm enforce` command. VxFS relocates qualifying files in the order in which it encounters them during its file system directory tree scan. When `tier1` volumes are fully occupied, VxFS stops scheduling qualifying files for relocation.

VxFS file placement policies are able to control file placement across any number of placement classes. The following example illustrates a rule for relocating files with low I/O temperatures from `tier1` volumes to `tier2` volumes, and to `tier3` volumes when `tier2` volumes are fully occupied:

```

<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS>tier1</CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
    <DESTINATION>
      <CLASS>tier3</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <IOTEMP Type="nrbytes">
      <MAX Flags="lt">4</MAX>
      <PERIOD>3</PERIOD>
    </IOTEMP>
  </WHEN>
</RELOCATE>

```

This rule relocates files whose 3-day I/O temperatures are less than 4 and which reside on `tier1` volumes. When VxFS calculates that already-relocated files would result in `tier2` volumes being fully occupied, VxFS relocates qualifying files to `tier3` volumes instead. VxFS relocates qualifying files as it encounters them in its scan of the file system directory tree.

The `<FROM>` clause in the `RELOCATE` statement is optional. If the clause is not present, VxFS evaluates files designated by the rule's `SELECT` statement for relocation no matter which volumes they reside on when the `fsppadm enforce` command is issued. The following example illustrates a fragment of a policy rule that relocates files according to their sizes, no matter where they reside when the `fsppadm enforce` command is issued:

```
<RELOCATE>
  <TO>
    <DESTINATION>
      <CLASS>tier1</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <SIZE Units="MB">
      <MAX Flags="lt">10</MAX>
    </SIZE>
  </WHEN>
</RELOCATE>
<RELOCATE>
  <TO>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <SIZE Units="MB">
      <MIN Flags="gteq">10</MIN>
      <MAX Flags="lt">100</MAX>
    </SIZE>
  </WHEN>
</RELOCATE>
<RELOCATE>
  <TO>
    <DESTINATION>
      <CLASS>tier3</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <SIZE Units="MB">
      <MIN Flags="gteq">100</MIN>
    </SIZE>
```

```
</WHEN>
</RELOCATE>
```

This rule relocates files smaller than 10 megabytes to `tier1` volumes, files between 10 and 100 megabytes to `tier2` volumes, and files larger than 100 megabytes to `tier3` volumes. VxFS relocates all qualifying files that do not already reside on volumes in their `DESTINATION` placement classes when the `fspadm enforce` command is issued.

DELETE statement

The `DELETE` file placement policy rule statement is very similar to the `RELOCATE` statement in both form and function, lacking only the `<TO>` clause. File placement policy-based deletion may be thought of as relocation with a fixed destination.

Note: Use `DELETE` statements with caution.

The following XML snippet illustrates the general form of the `DELETE` statement:

```
<DELETE>
  <FROM>
    <SOURCE>
      <CLASS> placement_class_name </CLASS>
    </SOURCE>
    <SOURCE>
      additional_placement_class_specifications
    </SOURCE>
  </FROM>
  <WHEN> relocation_conditions </WHEN>
</DELETE>
```

A `DELETE` statement contains the following clauses:

<code><FROM></code>	An optional clause that contains a list of placement classes from whose volumes designated files should be deleted if the files meet the conditions specified in the <code><WHEN></code> clause. No priority is associated with the ordering of placement classes in a <code><FROM></code> clause. If a file to which the rule applies is located on a volume in any specified placement class, the file is deleted. If a <code>DELETE</code> statement does not contain a <code><FROM></code> clause, VxFS deletes qualifying files no matter on which of a file system's volumes the files reside.

<WHEN> An optional clause specifying the conditions under which files to which the rule applies should be deleted. The form of the **<WHEN>** clause in a **DELETE** statement is identical to that of the **<WHEN>** clause in a **RELOCATE** statement. If a **DELETE** statement does not contain a **<WHEN>** clause, files designated by the rule's **SELECT** statement, and the **<FROM>** clause if it is present, are deleted unconditionally.

DELETE statement examples

The following example illustrates the use of the **DELETE** statement:

```
<DELETE>
  <FROM>
    <SOURCE>
      <CLASS>tier3</CLASS>
    </SOURCE>
  </FROM>
</DELETE>
<DELETE>
  <FROM>
    <SOURCE>
      <CLASS>tier2</CLASS>
    </SOURCE>
  </FROM>
  <WHEN>
    <ACCAGE Units="days">
      <MIN Flags="gt">120</MIN>
    </ACCAGE>
  </WHEN>
</DELETE>
```

The first **DELETE** statement unconditionally deletes files designated by the rule's **SELECT** statement that reside on **tier3** volumes when the **fsppadm enforce** command is issued. The absence of a **<WHEN>** clause in the **DELETE** statement indicates that deletion of designated files is unconditional.

The second **DELETE** statement deletes files to which the rule applies that reside on **tier2** volumes when the **fsppadm enforce** command is issued and that have not been accessed for the past 120 days.

Calculating I/O temperature and access temperature

An important application of VxFS SmartTier is automating the relocation of inactive files to lower cost storage. If a file has not been accessed for the period of time specified in the `<ACCAGE>` element, a scan of the file system should schedule the file for relocation to a lower tier of storage. But, time since last access is inadequate as the only criterion for activity-based relocation.

Why time since last access is inadequate as the only criterion for activity-based relocation:

- Access age is a binary measure. The time since last access of a file is computed by subtracting the time at which the `fsppadm enforce` command is issued from the POSIX `atime` in the file's metadata. If a file is opened the day before the `fsppadm enforce` command, its time since last access is one day, even though it may have been inactive for the month preceding. If the intent of a policy rule is to relocate inactive files to lower tier volumes, it will perform badly against files that happen to be accessed, however casually, within the interval defined by the value of the `<ACCAGE>` pa-rameter.
- Access age is a poor indicator of resumption of significant activity. Using `ACCAGE`, the time since last access, as a criterion for relocating inactive files to lower tier volumes may fail to schedule some relocations that should be performed, but at least this method results in less relocation activity than necessary. Using `ACCAGE` as a criterion for relocating previously inactive files that have become active is worse, because this method is likely to schedule relocation activity that is not warranted. If a policy rule's intent is to cause files that have experienced I/O activity in the recent past to be relocated to higher performing, perhaps more failure tolerant storage, `ACCAGE` is too coarse a filter. For example, in a rule specifying that files on `tier2` volumes that have been accessed within the last three days should be relocated to `tier1` volumes, no distinction is made between a file that was browsed by a single user and a file that actually was used intensively by applications.

SmartTier implements the concept of I/O temperature and access temperature to overcome these deficiencies. A file's I/O temperature is equal to the number of bytes transferred to or from it over a specified period of time divided by the size of the file. For example, if a file occupies one megabyte of storage at the time of an `fsppadm enforce` operation and the data in the file has been completely read or written 15 times within the last three days, VxFS calculates its 3-day average I/O temperature to be 5 (15 MB of I/O ÷ 1 MB file size ÷ 3 days).

Similarly, a file's average access temperature is the number of read or write requests made to it over a specified number of 24-hour periods divided by the number of periods. Unlike I/O temperature, access temperature is unrelated to file size. A large

file to which 20 I/O requests are made over a 2-day period has the same average access temperature as a small file accessed 20 times over a 2-day period.

If a file system's active placement policy includes any `<IOTEMP>` or `<ACCESSTEMP>` clauses, VxFS begins policy enforcement by using information in the file system's FCL file to calculate average I/O activity against all files in the file system during the longest `<PERIOD>` specified in the policy. Shorter specified periods are ignored. VxFS uses these calculations to qualify files for I/O temperature-based relocation and deletion.

See [“About the File Change Log file”](#) on page 931.

Note: If FCL is turned off, I/O temperature-based relocation will not be accurate. When you invoke the `fsppadm enforce` command, the command displays a warning if the FCL is turned off.

As its name implies, the File Change Log records information about changes made to files in a VxFS file system. In addition to recording creations, deletions, extensions, the FCL periodically captures the cumulative amount of I/O activity (number of bytes read and written) on a file-by-file basis. File I/O activity is recorded in the FCL each time a file is opened or closed, as well as at timed intervals to capture information about files that remain open for long periods.

If a file system's active file placement policy contains `<IOTEMP>` clauses, execution of the `fsppadm enforce` command begins with a scan of the FCL to extract I/O activity information over the period of interest for the policy. The period of interest is the interval between the time at which the `fsppadm enforce` command was issued and that time minus the largest interval value specified in any `<PERIOD>` element in the active policy.

For files with I/O activity during the largest interval, VxFS computes an approximation of the amount of read, write, and total data transfer (the sum of the two) activity by subtracting the I/O levels in the oldest FCL record that pertains to the file from those in the newest. It then computes each file's I/O temperature by dividing its I/O activity by its size at `Tscan`. Dividing by file size is an implicit acknowledgement that relocating larger files consumes more I/O resources than relocating smaller ones. Using this algorithm requires that larger files must have more activity against them in order to reach a given I/O temperature, and thereby justify the resource cost of relocation.

While this computation is an approximation in several ways, it represents an easy to compute, and more importantly, unbiased estimate of relative recent I/O activity upon which reasonable relocation decisions can be based.

File relocation and deletion decisions can be based on read, write, or total I/O activity.

The following XML snippet illustrates the use of `IOTEMP` in a policy rule to specify relocation of low activity files from `tier1` volumes to `tier2` volumes:

```
<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS>tier1</CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <IOTEMP Type="nrwbytes">
      <MAX Flags="lt">3</MAX>
      <PERIOD Units="days">4</PERIOD>
    </IOTEMP>
  </WHEN>
</RELOCATE>
```

This snippet specifies that files to which the rule applies should be relocated from `tier1` volumes to `tier2` volumes if their I/O temperatures fall below 3 over a period of 4 days. The `Type="nrwbytes"` XML attribute specifies that total data transfer activity, which is the the sum of bytes read and bytes written, should be used in the computation. For example, a 50 megabyte file that experienced less than 150 megabytes of data transfer over the 4-day period immediately preceding the `fsspadm enforce` scan would be a candidate for relocation. VxFS considers files that experience no activity over the period of interest to have an I/O temperature of zero. VxFS relocates qualifying files in the order in which it encounters the files in its scan of the file system directory tree.

Using I/O temperature or access temperature rather than a binary indication of activity, such as the POSIX `atime` or `mtime`, minimizes the chance of not relocating files that were only accessed occasionally during the period of interest. A large file that has had only a few bytes transferred to or from it would have a low I/O temperature, and would therefore be a candidate for relocation to `tier2` volumes, even if the activity was very recent.

But, the greater value of I/O temperature or access temperature as a file relocation criterion lies in upward relocation: detecting increasing levels of I/O activity against

files that had previously been relocated to lower tiers in a storage hierarchy due to inactivity or low temperatures, and relocating them to higher tiers in the storage hierarchy.

The following XML snippet illustrates relocating files from `tier2` volumes to `tier1` when the activity level against them increases.

```
<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS>tier2</CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
      <CLASS>tier1</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <IOTEMP Type="nrbytes">
      <MAX Flags="gt">5</MAX>
      <PERIOD Units="days">2</PERIOD>
    </IOTEMP>
  </WHEN>
</RELOCATE>
```

The `<RELOCATE>` statement specifies that files on `tier2` volumes whose I/O temperature as calculated using the number of bytes read is above 5 over a 2-day period are to be relocated to `tier1` volumes. Bytes written to the file during the period of interest are not part of this calculation.

Using I/O temperature rather than a binary indicator of activity as a criterion for file relocation gives administrators a granular level of control over automated file relocation that can be used to attune policies to application requirements. For example, specifying a large value in the `<PERIOD>` element of an upward relocation statement prevents files from being relocated unless I/O activity against them is sustained. Alternatively, specifying a high temperature and a short period tends to relocate files based on short-term intensity of I/O activity against them.

I/O temperature and access temperature utilize the `sqlite3` database for building a temporary table indexed on an inode. This temporary table is used to filter files based on I/O temperature and access temperature. The temporary table is stored in the database file `.__fspadm_fcliotemp.db`, which resides in the `lost+found` directory of the mount point.

Multiple criteria in file placement policy rule statements

In certain cases, file placement policy rule statements may contain multiple clauses that affect their behavior. In general, when a rule statement contains multiple clauses of a given type, all clauses must be satisfied in order for the statement to be effective. There are four cases of note in which multiple clauses may be used.

Multiple file selection criteria in SELECT statement clauses

Within a single `SELECT` statement, all the selection criteria clauses of a single type are treated as a selection list. A file need only satisfy a single criterion of a given type to be designated.

In the following example, files in any of the `db/datafiles`, `db/indexes`, and `db/logs` directories, all relative to the file system mount point, would be selected:

```
<SELECT>
  <DIRECTORY Flags="nonrecursive">db/datafiles</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/indexes</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/logs</DIRECTORY>
</SELECT>
```

This example is in direct contrast to the treatment of selection criteria clauses of different types. When a `SELECT` statement includes multiple types of file selection criteria, a file must satisfy one criterion of each type in order for the rule's action statements to apply.

In the following example, a file must reside in one of `db/datafiles`, `db/indexes`, or `db/logs` and be owned by one of `DBA_Manager`, `MFG_DBA`, or `HR_DBA` to be designated for possible action:

```
<SELECT>
  <DIRECTORY Flags="nonrecursive">db/datafiles</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/indexes</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/logs</DIRECTORY>
  <USER>DBA_Manager</USER>
  <USER>MFG_DBA</USER>
  <USER>HR_DBA</USER>
</SELECT>
```

If a rule includes multiple `SELECT` statements, a file need only satisfy one of them to be selected for action. This property can be used to specify alternative conditions for file selection.

In the following example, a file need only reside in one of db/datafiles, db/indexes, or db/logs or be owned by one of DBA_Manager, MFG_DBA, or HR_DBA to be designated for possible action:

```
<SELECT>
  <DIRECTORY Flags="nonrecursive">db/datafiles</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/indexes</DIRECTORY>
  <DIRECTORY Flags="nonrecursive">db/logs</DIRECTORY>
</SELECT>
<SELECT>
  <USER>DBA_Manager</USER>
  <USER>MFG_DBA</USER>
  <USER>HR_DBA</USER>
</SELECT>
```

Multiple placement classes in <ON> clauses of CREATE statements and in <TO> clauses of RELOCATE statements

Both the <ON> clause of the CREATE statement and the <TO> clause of the RELOCATE statement can specify priority ordered lists of placement classes using multiple <DESTINATION> XML elements. VxFS uses a volume in the first placement class in a list for the designated purpose of file creation or relocation, if possible. If no volume in the first listed class has sufficient free space or if the file system's volume set does not contain any volumes with that placement class, VxFS uses a volume in the second listed class if possible. If no volume in the second listed class can be used, a volume in the third listed class is used if possible, and so forth.

The following example illustrates of three placement classes specified in the <ON> clause of a CREATE statement:

```
<CREATE>
  <ON>
    <DESTINATION>
      <CLASS>tier1</CLASS>
    </DESTINATION>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
    <DESTINATION>
      <CLASS>tier3</CLASS>
    </DESTINATION>
  </ON>
</CREATE>
```

In this statement, VxFS would allocate space for newly created files designated by the rule's `SELECT` statement on `tier1` volumes if space was available. If no `tier1` volume had sufficient free space, VxFS would attempt to allocate space on a `tier2` volume. If no `tier2` volume had sufficient free space, VxFS would attempt allocation on a `tier3` volume. If sufficient space could not be allocated on a volume in any of the three specified placement classes, allocation would fail with an `ENOSPC` error, even if the file system's volume set included volumes in other placement classes that did have sufficient space.

The `<TO>` clause in the `RELOCATE` statement behaves similarly. VxFS relocates qualifying files to volumes in the first placement class specified if possible, to volumes in the second specified class if not, and so forth. If none of the destination criteria can be met, such as if all specified classes are fully occupied, qualifying files are not relocated, but no error is signaled in this case.

Multiple placement classes in `<FROM>` clauses of `RELOCATE` and `DELETE` statements

The `<FROM>` clause in `RELOCATE` and `DELETE` statements can include multiple source placement classes. However, unlike the `<ON>` and `<TO>` clauses, no order or priority is implied in `<FROM>` clauses. If a qualifying file resides on a volume in any of the placement classes specified in a `<FROM>` clause, it is relocated or deleted regardless of the position of its placement class in the `<FROM>` clause list of classes.

Multiple conditions in `<WHEN>` clauses of `RELOCATE` and `DELETE` statements

The `<WHEN>` clause in `RELOCATE` and `DELETE` statements may include multiple relocation criteria. Any or all of `<ACCAGE>`, `<MODAGE>`, `<SIZE>`, and `<IOTEMP>` can be specified. When multiple conditions are specified, all must be satisfied in order for a selected file to qualify for relocation or deletion.

In the following example, a selected file would have to be both inactive, that is, not accessed, for more than 30 days and larger than 100 megabytes to be eligible for relocation or deletion:

```
<WHEN>
  <ACCAGE Units="days">
    <MIN Flags="gt">30</MIN>
  </ACCAGE>
  <SIZE Units="MB">
    <MIN Flags="gt">100</MIN>
```

```
</SIZE>
</WHEN>
```

You cannot write rules to relocate or delete a single designated set of files if the files meet one of two or more relocation or deletion criteria.

File placement policy rule and statement ordering

You can use the SmartTier graphical user interface (GUI) to create any of four types of file placement policy documents. Alternatively, you can use a text editor or XML editor to create XML policy documents directly. The GUI places policy rule statements in the correct order to achieve the desired behavior. If you use a text editor, it is your responsibility to order policy rules and the statements in them so that the desired behavior results.

The rules that comprise a placement policy may occur in any order, but during both file allocation and `fsppadm enforce` relocation scans, the first rule in which a file is designated by a `SELECT` statement is the only rule against which that file is evaluated. Thus, rules whose purpose is to supersede a generally applicable behavior for a special class of files should precede the general rules in a file placement policy document.

The following XML snippet illustrates faulty rule placement with potentially unintended consequences:

```
<?xml version="1.0"?>
<!DOCTYPE FILE_PLACEMENT_POLICY SYSTEM "placement.dtd">
<FILE_PLACEMENT_POLICY Version="5.0">
  <RULE Name="GeneralRule">
    <SELECT>
      <PATTERN>*</PATTERN>
    </SELECT>
    <CREATE>
      <ON>
        <DESTINATION>
          <CLASS>tier2</CLASS>
        </DESTINATION>
      </ON>
    </CREATE>
    other_statements
  </RULE>
  <RULE Name="DatabaseRule">
    <SELECT>
      <PATTERN>*.db</PATTERN>
```

```

</SELECT>
<CREATE>
  <ON>
    <DESTINATION>
      <CLASS>tier1</CLASS>
    </DESTINATION>
  </ON>
</CREATE>
  other_statements
</RULE>
</FILE_PLACEMENT_POLICY>

```

The `GeneralRule` rule specifies that all files created in the file system, designated by `<PATTERN>*``</PATTERN>`, should be created on `tier2` volumes. The `DatabaseRule` rule specifies that files whose names include an extension of `.db` should be created on `tier1` volumes. The `GeneralRule` rule applies to any file created in the file system, including those with a naming pattern of `*.db`, so the `DatabaseRule` rule will never apply to any file. This fault can be remedied by exchanging the order of the two rules. If the `DatabaseRule` rule occurs first in the policy document, VxFS encounters it first when determining where to new place files whose names follow the pattern `*.db`, and correctly allocates space for them on `tier1` volumes. For files to which the `DatabaseRule` rule does not apply, VxFS continues scanning the policy and allocates space according to the specification in the `CREATE` statement of the `GeneralRule` rule.

A similar consideration applies to statements within a placement policy rule. VxFS processes these statements in order, and stops processing on behalf of a file when it encounters a statement that pertains to the file. This can result in unintended behavior.

The following XML snippet illustrates a `RELOCATE` statement and a `DELETE` statement in a rule that is intended to relocate if the files have not been accessed in 30 days, and delete the files if they have not been accessed in 90 days:

```

<RELOCATE>
  <TO>
    <DESTINATION>
      <CLASS>tier2</CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <ACCAGE Units="days">
      <MIN Flags="gt">30</MIN>
    </ACCAGE>

```

```
</WHEN>  
</RELOCATE>  
<DELETE>  
  <WHEN>  
    <ACCAGE Units="days">  
      <MIN Flags="gt">90</MIN>  
    </ACCAGE>  
  </WHEN>  
</DELETE>
```

As written with the `RELOCATE` statement preceding the `DELETE` statement, files will never be deleted, because the `<WHEN>` clause in the `RELOCATE` statement applies to all selected files that have not been accessed for at least 30 days. This includes those that have not been accessed for 90 days. VxFS ceases to process a file against a placement policy when it identifies a statement that applies to that file, so the `DELETE` statement would never occur. This example illustrates the general point that `RELOCATE` and `DELETE` statements that specify less inclusive criteria should precede statements that specify more inclusive criteria in a file placement policy document. The GUI automatically produce the correct statement order for the policies it creates.

File placement policies and extending files

In a VxFS file system with an active file placement policy, the placement class on whose volume a file resides is part of its metadata, and is attached when it is created and updated when it is relocated. When an application extends a file, VxFS allocates the incremental space on the volume occupied by the file if possible. If not possible, VxFS allocates the space on another volume in the same placement class. For example, if a file is created on a `tier1` volume and later relocated to a `tier2` volume, extensions to the file that occur before the relocation have space allocated on a `tier1` volume, while those occurring after to the relocation have their space allocated on `tier2` volumes. When a file is relocated, all of its allocated space, including the space acquired by extension, is relocated to `tier2` volumes in this case.

Using SmartTier with solid state disks

The SmartTier placement policies support SSD-based tiers with the following features:

- Allowance of fine grained temperatures, such as allowing hours as units for the `<IOTEMP>` and `<ACCESSTEMP>` criteria
See [“Fine grain temperatures with solid state disks”](#) on page 763.

- Support of the `Prefer` attribute for the `<IOTEMP>` and `<ACCESSTEMP>` criteria
See [“Prefer mechanism with solid state disks”](#) on page 764.
- Provision of a mechanism to relocate based on average I/O activity
See [“Average I/O activity with solid state disks”](#) on page 764.
- Reduction of the intensity and duration of scans to minimize the impact on resources, such as memory, CPU, and I/O bandwidth
See [“Frequent SmartTier scans with solid state disks”](#) on page 765.
- Quick identification of cold files
See [“Quick identification of cold files with solid state disks”](#) on page 765.

To gain these benefits, you must modify the existing placement policy as per the latest version of the DTD and assign the policy again. However, existing placement policies continue to function as before. You do not need to update the placement policies if you do not use the new features.

Fine grain temperatures with solid state disks

Before the solid state disk (SSD) enhancements, the SmartTier feature computed temperature values on a day granularity. Day granularity is the I/O activity per day over at least one day. As such, the `<PERIOD>` element had to be in days for the `<IOTEMP>` and `<ACCESSTEMP>` criteria. With SSDs, relocation decisions might need to happen within the day itself, based on I/O activity that Veritas File System (VxFS) measured over a shorter duration. As such, you can now specify "hours" for the `Units` attribute value for the `<IOTEMP>` and `<ACCESSTEMP>` criteria.

See [“Specifying the I/O temperature relocation criterion”](#) on page 743.

The following placement policy snippet gives an example of specifying 4 hours as the period of time:

```
<RELOCATE>
...
<WHEN>
  <IOTEMP Type="nwbytes">
    <MIN Flags="gteq"> 2 </MIN>
    <PERIOD Units="hours"> 4 </PERIOD>
  </IOTEMP>
</WHEN>
</RELOCATE>
```

Prefer mechanism with solid state disks

You can now specify a value for the `Prefer` attribute for the `<IOTEMP>` and `<ACCESSTEMP>` criteria, which gives preference to relocating files.

See [“Prefer attribute”](#) on page 744.

In case of a solid state disk (SSD)-based tier, you might want to relocate a file to an SSD as soon as there is a marked increase in the I/O activity. However, once Veritas File System (VxFS) has relocated the file to an SSD, it may be beneficial to keep the file on the SSD as long as the activity remains high to avoid frequent thrashing. You want to watch the activity for some time longer than the time that you watched the activity when you relocated the file to the SSD before you decide to move the file off of the SSD.

The following placement policy snippet gives an example of the `Prefer` criteria:

```
<RELOCATE>
...
<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high">
    <MIN Flags="gteq"> 3.4 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</RELOCATE>
```

If there are a number of files whose I/O temperature is greater than the given minimum value, the files with the higher temperature are first subject to the `RELOCATE` operation before the files with the lower temperature. This is particularly useful in case of SSDs, which are limited in size and are expensive. As such, you generally want to use SSDs for the most active files.

Average I/O activity with solid state disks

Before the solid state disk (SSD) enhancements, you were required to specify an absolute value of the temperature when you used the `ACCESSTEMP` criteria and `IOTEMP` criteria in the SmartTier placement policies. However, arriving at such absolute numbers is difficult and requires you to experiment and observe data access patterns over a period of time. Moreover, over a period of time, you might have to change this value due to changing access patterns. As such, you might need to repeat the experiment. To ease constructing `ACCESSTEMP` and `IOTEMP`-based policies, a new criteria has been introduced: `Average`.

See [“Average I/O activity criteria”](#) on page 745.

Frequent SmartTier scans with solid state disks

You can specify "hours" for the `Units` attribute value, and as such the I/O stats collection PERIOD can be much shorter than in previous releases. When not using solid state disks (SSDs), you can only specify "days" for the `Units` attribute value, which might be sufficient for your needs. However, a PERIOD shorter than a day is required in the context of using SSDs since the candidate files and their activity levels can change during the day. As a result, SmartTier must scan more frequently, which leads to a higher scan load on the host systems.

You must satisfy the following conflicting requirements simultaneously:

- Bring down the temperature collection windows to hourly levels.
- Reduce the impact of more frequent scans on resources, such as CPU, I/O, and memory.

The following scheme is an example of one way to reduce the impact of frequent scans:

- Confine the scan to only active files during the PERIOD by focusing only on the files that showed any activity in the File Change Log (FCL) by running the `fspadm` command with the `-C` option.
See [“Quick identification of cold files with solid state disks”](#) on page 765.
- Scan frequently, such as every few hours. Frequent scans potentially reduce the number of inodes that VxFS touches and logs in the File Change Log (FCL) file, thereby limiting the duration of each scan. As such, the changes that VxFS collects in the FCL file since the last scan provide details on fewer active files.
- Use the `<IOTEMP>` and `<ACCESSTEMP>` criteria to promote files to SSDs more aggressively, which leaves cold files sitting in SSDs.

Quick identification of cold files with solid state disks

The placement mechanism generally leaves the cold files in solid state disks (SSDs) if the files continue to remain inactive. This results in a lack of room for active files if the active files need to be moved into SSDs, and thus results in ineffective use of storage. An SSD enhancement for identifying cold files quickly solves this problem.

The enhancement is a method for quickly identifying files on a particular tier of the SmartTier file system so that the files can be relocated if necessary. The method consists of a map that associates storage devices with the inodes of files residing on the storage devices.

Veritas File System (VxFS) updates the file location map during the following times:

- SmartTier's own file relocations

- On examination of the file system's File Change Log (FCL) for changes that are made outside of SmartTier's scope.

Both of these updates occur during SmartTier's relocation scans, which are typically scheduled to occur periodically. But, you can also update the file location map anytime by running the `fsppadm` command with the `-T` option.

The `-C` option is useful to process active files before any other files. For best results, specify the `-T` option in conjunction with the `-C` option. Specifying both the `-T` option and `-C` option causes the `fsppadm` command to evacuate any cold files first to create room in the SSD tier to accommodate any active files that will be moved into the SSD tier via the `-C` option. Specifying `-C` in conjunction with `-T` confines the scope of the scan, which consumes less time and resources, and thus allows frequent scans to meet the dynamic needs of data placement.

See [“Enforcing a placement policy”](#) on page 729.

See the `fsppadm(1M)` manual page.

With the help of the map, instead of scanning the full file system, you can confine the scan to only the files on the SSD tiers in addition to the active files that VxFS recorded in the FCL. This scheme potentially achieves the dual purpose of reducing the temperature time granularity and at the same time reducing the scan load.

Example placement policy when using solid state disks

The following snippet is one possible placement policy for use with solid state disk (SSD)-based tiers.

```
<?xml version="1.0"?>
<!DOCTYPE PLACEMENT_POLICY SYSTEM "/opt/VRTSvxfs/etc/placement_policy.dtd">
<PLACEMENT_POLICY Version="5.0" Name="SSD_policy">
  <RULE Flags="data" Name="all_files">
    <COMMENT>
      The first two RELOCATES will do the evacuation
      out of SSDs to create room for any relocations
      into the SSDs by the third RELOCATE. The parameters
      that can be tuned are basically values for PERIOD and
      the values of MIN and/or MAX as the per the case.
      The values for MIN and MAX are treated as multiples of
      average activity over past 24 hour period.
    </COMMENT>
    <SELECT>
      <PATTERN> * </PATTERN>
    </SELECT>
```

```

<CREATE>
  <COMMENT>
    create files on ssdtier, failing which
    create them on other tiers
  </COMMENT>
  <ON>
    <DESTINATION Flags="any">
      <CLASS> ssdtier </CLASS>
    </DESTINATION>
  </ON>
</CREATE>

<RELOCATE>
  <COMMENT>
    Move the files out of SSD if their last 3 hour
    write IOTEMP is more than 1.5 times the last
    24 hour average write IOTEMP. The PERIOD is
    purposely shorter than the other RELOCATES
    because we want to move it out as soon as
    write activity starts peaking. This criteria
    could be used to reduce SSD wear outs.
  </COMMENT>
  <FROM>
    <SOURCE>
      <CLASS> ssdtier </CLASS>
    </SOURCE>
  </FROM>
  <TO>
    <DESTINATION>
      <CLASS> nonssd_tier </CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <IOTEMP Type="nwbytes" Average="*">
      <MIN Flags="gt"> 1.5 </MIN>
      <PERIOD Units="hours"> 3 </PERIOD>
    </IOTEMP>
  </WHEN>
</RELOCATE>

<RELOCATE>
  <COMMENT>
    OR move the files out of SSD if their last 6 hour

```

```

        read IOTEMP is less than half the last 24 hour
        average read IOTEMP. The PERIOD is longer,
        we may want to observe longer periods
        having brought the file in. This avoids quickly
        sending the file out of SSDs once in.
    </COMMENT>
    <FROM>
        <SOURCE>
            <CLASS> ssdtier </CLASS>
        </SOURCE>
    </FROM>
    <TO>
        <DESTINATION>
            <CLASS> nonssd_tier </CLASS>
        </DESTINATION>
    </TO>
    <WHEN>
        <IOTEMP Type="nrbytes" Average="*">
            <MAX Flags="lt"> 0.5 </MAX>
            <PERIOD Units="hours"> 6 </PERIOD>
        </IOTEMP>
    </WHEN>
</RELOCATE>

<RELOCATE>
    <COMMENT>
        OR move the files into SSD if their last 3 hour
        read IOTEMP is more than or equal to 1.5 times
        the last 24 hour average read IOTEMP AND
        their last 6 hour write IOTEMP is less than
        half of the last 24 hour average write IOTEMP
    </COMMENT>
    <TO>
        <DESTINATION>
            <CLASS> ssd_tier </CLASS>
        </DESTINATION>
    </TO>
    <WHEN>
        <IOTEMP Type="nrbytes" Prefer="high" Average="*">
            <MIN Flags="gteq"> 1.5 </MIN>
            <PERIOD Units="hours"> 3 </PERIOD>
        </IOTEMP>
        <IOTEMP Type="nwbytes" Average="*">


```

```

        <MAX Flags="lt"> 0.5 </MAX>
        <PERIOD Units="hours"> 3 </PERIOD>
    </IOTEMP>
</WHEN>
</RELOCATE>
</RULE>
</PLACEMENT_POLICY>
    
```

In this placement policy, new files are created on the SSD tiers if space is available, or elsewhere if space is not available. When enforce is performed, the files that are currently in SSDs whose write activity is increased above a threshold or whose read activity fell below a threshold over a given period are moved out of the SSDs. The first two RELOCATES capture this intent. However, the files whose read activity intensified above a threshold and whose write activity does not exceed a threshold over the given period are moved into SSDs, while giving preference to files with higher read activity.

The following figure illustrates the behavior of the example placement policy:



The files whose I/O activity falls in the light gray area are good candidates for moving in to SSD storage. These files have less write activity such that they have less impact on wear leveling, and the slower write times to SSDs is less of a factor. These files have intense read activity, which also makes the files ideal for placement on SSDs since read activity does not cause any wear leveling side effects, and

reads are faster from SSDs. In contrast, the files whose I/O activity falls in the dark gray area are good candidates to be moved out of SSD storage, since they have more write activity or less read activity. Greater write activity leads to greater wear leveling of the SSDs, and your file system's performance suffers from the slower write times of SSDs. Lesser read activity means that you are not benefitting from the faster read times of SSDs with these files.

Administering hot-relocation

This chapter includes the following topics:

- [About hot-relocation](#)
- [How hot-relocation works](#)
- [Configuring a system for hot-relocation](#)
- [Displaying spare disk information](#)
- [Marking a disk as a hot-relocation spare](#)
- [Removing a disk from use as a hot-relocation spare](#)
- [Excluding a disk from hot-relocation use](#)
- [Making a disk available for hot-relocation use](#)
- [Configuring hot-relocation to use only spare disks](#)
- [Moving relocated subdisks](#)
- [Modifying the behavior of hot-relocation](#)

About hot-relocation

If a volume has a disk I/O failure (for example, the disk has an uncorrectable error), Veritas Volume Manager (VxVM) can detach the plex involved in the failure. I/O stops on that plex but continues on the remaining plexes of the volume.

If a disk fails completely, VxVM can detach the disk from its disk group. All plexes on the disk are disabled. If there are any unmirrored volumes on a disk when it is detached, those volumes are also disabled.

Apparent disk failure may not be due to a fault in the physical disk media or the disk controller, but may instead be caused by a fault in an intermediate or ancillary component such as a cable, host bus adapter, or power supply.

The hot-relocation feature in VxVM automatically detects disk failures, and notifies the system administrator and other nominated users of the failures by electronic mail. Hot-relocation also attempts to use spare disks and free disk space to restore redundancy and to preserve access to mirrored and RAID-5 volumes.

See [“How hot-relocation works”](#) on page 772.

If hot-relocation is disabled or you miss the electronic mail, you can use the `vxprint` command or the graphical user interface to examine the status of the disks. You may also see driver error messages on the console or in the system messages file.

Failed disks must be removed and replaced manually.

See [“Removing and replacing disks”](#) on page 858.

For more information about recovering volumes and their data after hardware failure, see the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide*.

How hot-relocation works

Hot-relocation allows a system to react automatically to I/O failures on redundant (mirrored or RAID-5) VxVM objects, and to restore redundancy and access to those objects. VxVM detects I/O failures on objects and relocates the affected subdisks to disks designated as spare disks or to free space within the disk group. VxVM then reconstructs the objects that existed before the failure and makes them redundant and accessible again.

When a partial disk failure occurs (that is, a failure affecting only some subdisks on a disk), redundant data on the failed portion of the disk is relocated. Existing volumes on the unaffected portions of the disk remain accessible.

Hot-relocation is only performed for redundant (mirrored or RAID-5) subdisks on a failed disk. Non-redundant subdisks on a failed disk are not relocated, but the system administrator is notified of their failure.

Hot-relocation is enabled by default and takes effect without the intervention of the system administrator when a failure occurs.

The hot-relocation daemon, `vxrelocd`, detects and reacts to VxVM events that signify the following types of failures:

Disk failure	This is normally detected as a result of an I/O failure from a VxVM object. VxVM attempts to correct the error. If the error cannot be corrected, VxVM tries to access configuration information in the private region of the disk. If it cannot access the private region, it considers the disk failed.
Plex failure	This is normally detected as a result of an uncorrectable I/O error in the plex (which affects subdisks within the plex). For mirrored volumes, the plex is detached.
RAID-5 subdisk failure	This is normally detected as a result of an uncorrectable I/O error. The subdisk is detached.

When `vxrelocd` detects such a failure, it performs the following steps:

- `vxrelocd` informs the system administrator (and other nominated users) by electronic mail of the failure and which VxVM objects are affected.
See [“Partial disk failure mail messages”](#) on page 775.
See [“Complete disk failure mail messages”](#) on page 776.
See [“Modifying the behavior of hot-relocation”](#) on page 786.
- `vxrelocd` next determines if any subdisks can be relocated. `vxrelocd` looks for suitable space on disks that have been reserved as hot-relocation spares (marked `spare`) in the disk group where the failure occurred. It then relocates the subdisks to use this space.
- If no spare disks are available or additional space is needed, `vxrelocd` uses free space on disks in the same disk group, except those disks that have been excluded for hot-relocation use (marked `nohotuse`). When `vxrelocd` has relocated the subdisks, it reattaches each relocated subdisk to its plex.
- Finally, `vxrelocd` initiates appropriate recovery procedures. For example, recovery includes mirror resynchronization for mirrored volumes or data recovery for RAID-5 volumes. It also notifies the system administrator of the hot-relocation and recovery actions that have been taken.

If relocation is not possible, `vxrelocd` notifies the system administrator and takes no further action.

Warning: Hot-relocation does not guarantee the same layout of data or the same performance after relocation. An administrator should check whether any configuration changes are required after hot-relocation occurs.


Relocation of failing subdisks is not possible in the following cases:

- The failing subdisks are on non-redundant volumes (that is, volumes of types other than mirrored or RAID-5).
- There are insufficient spare disks or free disk space in the disk group.
- The only available space is on a disk that already contains a mirror of the failing plex.
- The only available space is on a disk that already contains the RAID-5 log plex or one of its healthy subdisks. Failing subdisks in the RAID-5 plex cannot be relocated.
- If a mirrored volume has a dirty region logging (DRL) log subdisk as part of its data plex, failing subdisks belonging to that plex cannot be relocated.
- If a RAID-5 volume log plex or a mirrored volume DRL log plex fails, a new log plex is created elsewhere. There is no need to relocate the failed subdisks of the log plex.

See the `vxrelocd(1M)` manual page.

[Figure 34-1](#) shows the hot-relocation process in the case of the failure of a single subdisk of a RAID-5 volume.

Figure 34-1 Example of hot-relocation for a subdisk in a RAID-5 volume



Partial disk failure mail messages

If hot-relocation is enabled when a plex or disk is detached by a failure, mail indicating the failed objects is sent to `root`. If a partial disk failure occurs, the mail identifies the failed plexes. For example, if a disk containing mirrored volumes fails, you can receive mail information as shown in the following example:

```
To: root
Subject: Volume Manager failures on host teal
Failures have been detected by the Veritas Volume Manager:

failed plexes:
home-02
src-02
```

Mail can be sent to users other than `root`.

See [“Modifying the behavior of hot-relocation”](#) on page 786.

You can determine which disk is causing the failures in the above example message by using the following command:

```
# vxstat -g mydg -s -ff home-02 src-02
```

The `-s` option asks for information about individual subdisks, and the `-ff` option displays the number of failed read and write operations. The following output display is typical:

TYP	NAME	FAILED	
		READS	WRITES
sd	mydg01-04	0	0
sd	mydg01-06	0	0
sd	mydg02-03	1	0
sd	mydg02-04	1	0

This example shows failures on reading from subdisks `mydg02-03` and `mydg02-04` of disk `mydg02`.

Hot-relocation automatically relocates the affected subdisks and initiates any necessary recovery procedures. However, if relocation is not possible or the hot-relocation feature is disabled, you must investigate the problem and attempt to recover the plexes. Errors can be caused by cabling failures, so check the cables connecting your disks to your system. If there are obvious problems, correct them and recover the plexes using the following command:

```
# vxrecover -b -g mydg home src
```

This starts recovery of the failed plexes in the background (the command prompt reappears before the operation completes). If an error message appears later, or if the plexes become detached again and there are no obvious cabling failures, replace the disk.

See [“Removing and replacing disks”](#) on page 858.

Complete disk failure mail messages

If a disk fails completely and hot-relocation is enabled, the mail message lists the disk that failed and all plexes that use the disk. For example, you can receive mail as shown in this example display:

```
To: root
Subject: Volume Manager failures on host teal
```

```
Failures have been detected by the Veritas Volume Manager:
```

```
failed disks:  
mydg02
```

```
failed plexes:  
home-02  
src-02  
mkting-01
```

```
failing disks:  
mydg02
```

This message shows that `mydg02` was detached by a failure. When a disk is detached, I/O cannot get to that disk. The plexes `home-02`, `src-02`, and `mkting-01` were also detached (probably because of the failure of the disk).

One possible cause of the problem could be a cabling error.

See [“Partial disk failure mail messages”](#) on page 775.

If the problem is not a cabling error, replace the disk.

See [“Removing and replacing disks”](#) on page 858.

How space is chosen for relocation

A spare disk must be initialized and placed in a disk group as a spare before it can be used for replacement purposes. If no disks have been designated as spares when a failure occurs, VxVM automatically uses any available free space in the disk group in which the failure occurs. If there is not enough spare disk space, a combination of spare space and free space is used.

When selecting space for relocation, hot-relocation preserves the redundancy characteristics of the VxVM object to which the relocated subdisk belongs. For example, hot-relocation ensures that subdisks from a failed plex are not relocated to a disk containing a mirror of the failed plex. If redundancy cannot be preserved using any available spare disks and/or free space, hot-relocation does not take place. If relocation is not possible, the system administrator is notified and no further action is taken.

From the eligible disks, hot-relocation attempts to use the disk that is “closest” to the failed disk. The value of “closeness” depends on the controller, target, and disk number of the failed disk. A disk on the same controller as the failed disk is closer than a disk on a different controller. A disk under the same target as the failed disk is closer than one on a different target.

Hot-relocation tries to move all subdisks from a failing drive to the same destination disk, if possible.

If the failing disk is a root disk, hot-relocation only works if it can relocate all of the file systems to the same disk. If none are found, the system administrator is notified through email.

When hot-relocation takes place, the failed subdisk is removed from the configuration database, and VxVM ensures that the disk space used by the failed subdisk is not recycled as free space.

Configuring a system for hot-relocation

By designating spare disks and making free space on disks available for use by hot relocation, you can control how disk space is used for relocating subdisks in the event of a disk failure. If the combined free space and space on spare disks is not sufficient or does not meet the redundancy constraints, the subdisks are not relocated.

Find out which disks are spares or are excluded from hot-relocation.

See [“Displaying spare disk information”](#) on page 779.

You can prepare for hot-relocation by designating one or more disks per disk group as hot-relocation spares.

See [“Marking a disk as a hot-relocation spare”](#) on page 779.

If required, you can remove a disk from use as a hot-relocation spare

See [“Removing a disk from use as a hot-relocation spare”](#) on page 780.

If no spares are available at the time of a failure or if there is not enough space on the spares, free space on disks in the same disk group as where the failure occurred is automatically used, unless it has been excluded from hot-relocation use.

See [“Excluding a disk from hot-relocation use”](#) on page 781.

See [“Making a disk available for hot-relocation use”](#) on page 782.

Depending on the locations of the relocated subdisks, you can choose to move them elsewhere after hot-relocation occurs.

See [“Configuring hot-relocation to use only spare disks”](#) on page 782.

After a successful relocation, remove and replace the failed disk.

See [“Removing and replacing disks”](#) on page 858.

Displaying spare disk information

Use the following command to display information about spare disks that are available for relocation:

```
# vxdg [-g diskgroup] spare
```

The following is example output:

GROUP	DISK	DEVICE	TAG	OFFSET	LENGTH	FLAGS
mydg	mydg02	c0t2d0s2	c0t2d0s2	0	658007	s

Here `mydg02` is the only disk designated as a spare in the `mydg` disk group. The `LENGTH` field indicates how much spare space is currently available on `mydg02` for relocation.

The following commands can also be used to display information about disks that are currently designated as spares:

- `vxdisk list` lists disk information and displays spare disks with a `spare` flag.
- `vxprint` lists disk and other information and displays spare disks with a `SPARE` flag.
- The `list` menu item on the `vxdiskadm` main menu lists all disks including spare disks.

Marking a disk as a hot-relocation spare

Hot-relocation allows the system to react automatically to I/O failure by relocating redundant subdisks to other disks. Hot-relocation then restores the affected VxVM objects and data. If a disk has already been designated as a spare in the disk group, the subdisks from the failed disk are relocated to the spare disk. Otherwise, any suitable free space in the disk group is used except for the free space on the disks that were previously excluded from hot-relocation use.

To designate a disk as a hot-relocation spare, enter the following command:

```
# vxedit [-g diskgroup] set spare=on diskname
```

where *diskname* is the disk media name.

For example, to designate `mydg01` as a spare in the disk group, `mydg`, enter the following command:

```
# vxedit -g mydg set spare=on mydg01
```

You can use the `vxdisk list` command to confirm that this disk is now a spare; `mydg01` should be listed with a `spare` flag.

Any VM disk in this disk group can now use this disk as a spare in the event of a failure. If a disk fails, hot-relocation automatically occurs (if possible). You are notified of the failure and relocation through electronic mail. After successful relocation, you may want to replace the failed disk.

To use `vxdiskadm` to designate a disk as a hot-relocation spare

- 1 Select **Mark a disk as a spare for a disk group** from the `vxdiskadm` main menu.
- 2 At the following prompt, enter a disk media name (such as `mydg01`):

```
Enter disk name [<disk>,list,q,?] mydg01
```

The following notice is displayed when the disk has been marked as spare:

```
VxVM NOTICE V-5-2-219 Marking of mydg01 in mydg as a spare disk
is complete.
```

- 3 At the following prompt, indicate whether you want to add more disks as spares (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Mark another disk as a spare? [y,n,q,?] (default: n)
```

Any VM disk in this disk group can now use this disk as a spare in the event of a failure. If a disk fails, hot-relocation should automatically occur (if possible). You should be notified of the failure and relocation through electronic mail. After successful relocation, you may want to replace the failed disk.

Removing a disk from use as a hot-relocation spare

While a disk is designated as a spare, the space on that disk is not used for the creation of VxVM objects within its disk group. If necessary, you can free a spare disk for general use by removing it from the pool of hot-relocation disks.

To remove a spare from the hot-relocation pool, use the following command:

```
# vxedit [-g diskgroup] set spare=off diskname
```

where *diskname* is the disk media name.

For example, to make `mydg01` available for normal use in the disk group, `mydg`, use the following command:

```
# vxedit -g mydg set spare=off mydg01
```


To use vxdiskadm to remove a disk from the hot-relocation pool

- 1 Select **Turn off the spare flag on a disk** from the vxdiskadm main menu.
- 2 At the following prompt, enter the disk media name of a spare disk (such as mydg01):

Enter disk name [<disk>,list,q,?] **mydg01**

The following confirmation is displayed:

VxVM NOTICE V-5-2-143 Disk mydg01 in mydg no longer marked as a spare disk.

- 3 At the following prompt, indicate whether you want to disable more spare disks (y) or return to the vxdiskadm main menu (n):

Turn off spare flag on another disk? [y,n,q,?] (default: n)

Excluding a disk from hot-relocation use

To exclude a disk from hot-relocation use, use the following command:

```
# vxedit [-g diskgroup] set nohotuse=on diskname
```

where *diskname* is the disk media name.

To use vxdiskadm to exclude a disk from hot-relocation use

- 1 Select **Exclude a disk from hot-relocation use** from the vxdiskadm main menu.
- 2 At the following prompt, enter the disk media name (such as mydg01):

Enter disk name [<disk>,list,q,?] **mydg01**

The following confirmation is displayed:

VxVM INFO V-5-2-925 Excluding mydg01 in mydg from hot-relocation use is complete.

- 3 At the following prompt, indicate whether you want to add more disks to be excluded from hot-relocation (y) or return to the vxdiskadm main menu (n):

Exclude another disk from hot-relocation use? [y,n,q,?] (default: n)

Making a disk available for hot-relocation use

Free space is used automatically by hot-relocation in case spare space is not sufficient to relocate failed subdisks. You can limit this free space usage by hot-relocation by specifying which free disks should not be touched by hot-relocation. If a disk was previously excluded from hot-relocation use, you can undo the exclusion and add the disk back to the hot-relocation pool.

To make a disk available for hot-relocation use, use the following command:

```
# vxedit [-g diskgroup] set nohotuse=off diskname
```

To use `vxdiskadm` to make a disk available for hot-relocation use

- 1 Select `Make a disk available for hot-relocation use` from the `vxdiskadm` main menu.
- 2 At the following prompt, enter the disk media name (such as `mydg01`):

```
Enter disk name [<disk>,list,q,?] mydg01
```

The following confirmation is displayed:

```
V-5-2-932 Making mydg01 in mydg available for hot-relocation  
use is complete.
```

- 3 At the following prompt, indicate whether you want to add more disks to be excluded from hot-relocation (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Make another disk available for hot-relocation use? [y,n,q,?]  
(default: n)
```

Configuring hot-relocation to use only spare disks

If you want VxVM to use only spare disks for hot-relocation, add the following line to the file `/etc/default/vxassist`:

```
spare=only
```

If not enough storage can be located on disks marked as spare, the relocation fails. Any free space on non-spare disks is not used.

Moving relocated subdisks

When hot-relocation occurs, subdisks are relocated to spare disks and/or available free space within the disk group. The new subdisk locations may not provide the same performance or data layout that existed before hot-relocation took place. You can move the relocated subdisks (after hot-relocation is complete) to improve performance.

You can also move the relocated subdisks of the spare disks to keep the spare disk space free for future hot-relocation needs. Another reason for moving subdisks is to recreate the configuration that existed before hot-relocation occurred.

During hot-relocation, one of the electronic mail messages sent to `root` is shown in the following example:

```
To: root
Subject: Volume Manager failures on host teal

Attempting to relocate subdisk mydg02-03 from plex home-02.
Dev_offset 0 length 1164 dm_name mydg02 da_name c0t5d0s2.
The available plex home-01 will be used to recover the data.
```

This message has information about the subdisk before relocation and can be used to decide where to move the subdisk after relocation.

Here is an example message that shows the new location for the relocated subdisk:

```
To: root
Subject: Attempting VxVM relocation on host teal

Volume home Subdisk mydg02-03 relocated to mydg05-01,
but not yet recovered.
```

Before you move any relocated subdisks, fix or replace the disk that failed.

See [“Removing and replacing disks”](#) on page 858.

Once this is done, you can move a relocated subdisk back to the original disk as described in the following sections.

Warning: During subdisk move operations, RAID-5 volumes are not redundant.

Moving relocated subdisks using vxunreloc

VxVM hot-relocation allows the system to automatically react to I/O failures on a redundant VxVM object at the subdisk level and then take necessary action to make

the object available again. This mechanism detects I/O failures in a subdisk, relocates the subdisk, and recovers the plex associated with the subdisk. After the disk has been replaced, `vxunreloc` allows you to restore the system back to the configuration that existed before the disk failure. `vxunreloc` allows you to move the hot-relocated subdisks back onto a disk that was replaced due to a failure.

When `vxunreloc` is invoked, you must specify the disk media name where the hot-relocated subdisks originally resided. When `vxunreloc` moves the subdisks, it moves them to the original offsets. If you try to unrellocate to a disk that is smaller than the original disk that failed, `vxunreloc` does nothing except return an error.

`vxunreloc` provides an option to move the subdisks to a different disk from where they were originally relocated. It also provides an option to unrellocate subdisks to a different offset as long as the destination disk is large enough to accommodate all the subdisks.

If `vxunreloc` cannot replace the subdisks back to the same original offsets, a force option is available that allows you to move the subdisks to a specified disk without using the original offsets.

See the `vxunreloc(1M)` manual page.

The examples in the following sections demonstrate the use of `vxunreloc`.

Moving hot-relocated subdisks back to their original disk

Assume that `mydg01` failed and all the subdisks were relocated. After `mydg01` is replaced, `vxunreloc` can be used to move all the hot-relocated subdisks back to `mydg01`.

```
# vxunreloc -g mydg mydg01
```

Moving hot-relocated subdisks back to a different disk

The `vxunreloc` utility provides the `-n` option to move the subdisks to a different disk from where they were originally relocated.

Assume that `mydg01` failed, and that all of the subdisks that resided on it were hot-relocated to other disks. `vxunreloc` provides an option to move the subdisks to a different disk from where they were originally relocated. After the disk is repaired, it is added back to the disk group using a different name, for example, `mydg05`. If you want to move all the hot-relocated subdisks back to the new disk, the following command can be used:

```
# vxunreloc -g mydg -n mydg05 mydg01
```

The destination disk should have at least as much storage capacity as was in use on the original disk. If there is not enough space, the `unrelocate` operation will fail and none of the subdisks will be moved.

Forcing hot-relocated subdisks to accept different offsets

By default, `vxunreloc` attempts to move hot-relocated subdisks to their original offsets. However, `vxunreloc` fails if any subdisks already occupy part or all of the area on the destination disk. In such a case, you have two choices:

- Move the existing subdisks somewhere else, and then re-run `vxunreloc`.
- Use the `-f` option provided by `vxunreloc` to move the subdisks to the destination disk, but leave it to `vxunreloc` to find the space on the disk. As long as the destination disk is large enough so that the region of the disk for storing subdisks can accommodate all subdisks, all the hot-relocated subdisks will be unrelocated without using the original offsets.

Assume that `mydg01` failed and the subdisks were relocated and that you want to move the hot-relocated subdisks to `mydg05` where some subdisks already reside. You can use the force option to move the hot-relocated subdisks to `mydg05`, but not to the exact offsets:

```
# vxunreloc -g mydg -f -n mydg05 mydg01
```

Examining which subdisks were hot-relocated from a disk

If a subdisk was hot relocated more than once due to multiple disk failures, it can still be unrelocated back to its original location. For instance, if `mydg01` failed and a subdisk named `mydg01-01` was moved to `mydg02`, and then `mydg02` experienced disk failure, all of the subdisks residing on it, including the one which was hot-relocated to it, will be moved again. When `mydg02` was replaced, a `vxunreloc` operation for `mydg02` will do nothing to the hot-relocated subdisk `mydg01-01`. However, a replacement of `mydg01` followed by a `vxunreloc` operation, moves `mydg01-01` back to `mydg01` if `vxunreloc` is run immediately after the replacement.

After the disk that experienced the failure is fixed or replaced, `vxunreloc` can be used to move all the hot-relocated subdisks back to the disk. When a subdisk is hot-relocated, its original disk-media name and the offset into the disk are saved in the configuration database. When a subdisk is moved back to the original disk or to a new disk using `vxunreloc`, the information is erased. The original disk-media name and the original offset are saved in the subdisk records. To print all of the subdisks that were hot-relocated from `mydg01` in the `mydg` disk group, use the following command:

```
# vxprint -g mydg -se 'sd_orig_dmname="mydg01"'
```

Restarting vxunreloc after errors

`vxunreloc` moves subdisks in three phases:

- `vxunreloc` creates as many subdisks on the specified destination disk as there are subdisks to be unrelocated. The string `UNRELOC` is placed in the `comment` field of each subdisk record.
Creating the subdisk is an all-or-nothing operation. If `vxunreloc` cannot create all the subdisks successfully, none are created, and `vxunreloc` exits.
- `vxunreloc` moves the data from each subdisk to the corresponding newly created subdisk on the destination disk.
- When all subdisk data moves have been completed successfully, `vxunreloc` sets the `comment` field to the null string for each subdisk on the destination disk whose `comment` field is currently set to `UNRELOC`.

The `comment` fields of all the subdisks on the destination disk remain marked as `UNRELOC` until phase 3 completes. If its execution is interrupted, `vxunreloc` can subsequently re-use subdisks that it created on the destination disk during a previous execution, but it does not use any data that was moved to the destination disk.

If a subdisk data move fails, `vxunreloc` displays an error message and exits. Determine the problem that caused the move to fail, and fix it before re-executing `vxunreloc`.

If the system goes down after the new subdisks are created on the destination disk, but before all the data has been moved, re-execute `vxunreloc` when the system has been rebooted.

Warning: Do not modify the string `UNRELOC` in the comment field of a subdisk record.

Modifying the behavior of hot-relocation

Hot-relocation is turned on as long as the `vxrelocd` process is running. You should normally leave hot-relocation turned on so that you can take advantage of this feature if a failure occurs. However, if you choose to disable hot-relocation (perhaps because you do not want the free space on your disks to be used for relocation), you can prevent `vxrelocd` from starting at system startup time by editing the startup file that invokes `vxrelocd`. This file is `/lib/svc/method/vxvm-recover` in Solaris 10, or `/etc/init.d/vxvm-recover` in previous releases of the Solaris OS.

If the hot-relocation daemon is disabled, then automatic storage reclamation on deleted volumes is also disabled.

You can alter the behavior of `vxrelocd` as follows:

- 1 To prevent `vxrelocd` starting, comment out the entry that invokes it in the startup file:

```
# nohup vxrelocd root &
```

- 2 By default, `vxrelocd` sends electronic mail to `root` when failures are detected and relocation actions are performed. You can instruct `vxrelocd` to notify additional users by adding the appropriate user names as shown here:

```
# nohup vxrelocd root user1 user2 &
```

- 3 To reduce the impact of recovery on system performance, you can instruct `vxrelocd` to increase the delay between the recovery of each region of the volume, as shown in the following example:

```
# nohup vxrelocd -o slow[=IOdelay] root &
```

where the optional *IOdelay* value indicates the desired delay in milliseconds. The default value for the delay is 250 milliseconds.

On a Solaris 10 system, after making changes to the way `vxrelocd` is invoked in the startup file, run the following command to notify that the service configuration has changed:

```
# svcadm refresh vxvm/vxvm-recover
```

See the `vxrelocd(1M)` manual page.

Deduplicating data on Solaris SPARC

This chapter includes the following topics:

- [About deduplicating data](#)
- [Deduplicating data](#)
- [Deduplication results](#)
- [Deduplication supportability](#)
- [Deduplication use cases](#)
- [Deduplication limitations](#)

About deduplicating data

The data deduplication feature eliminates duplicate blocks used by your data by comparing blocks across the file system. When the data deduplication feature finds a duplicate block, it removes the space used and instead creates a pointer to the common block. If you change the duplicate file, thus making the files no longer share the same block, then that changed block is saved to disk instead of the pointer. You can perform post-process periodic deduplication in a file system to eliminate duplicate data without any continuous cost in CPU overhead. You can verify whether data is duplicated on demand, and then efficiently and securely eliminate the duplicates. The deduplication process performs the following tasks:

- Scans the file system for changes
- Fingerprints the data
- Identifies duplicates

- Eliminates duplicates after verifying the duplicates

The amount of space savings that you get from deduplicating depends on your data. Deduplicating different data gives different space savings.

You deduplicate data using the `fsdedupadm` command.

See the `fsdedupadm(1M)` manual page.

Deduplication requires an Enterprise license.

About deduplication chunk size

The deduplication chunk size, which is also referred to as deduplication granularity, is the unit at which fingerprints are computed. A valid chunk size is between 4k and 128k and power of two. Once set, the only way to change the chunk size is to remove and re-enable deduplication on the file system.

You should carefully select the chunk size, as the size has significant impact on deduplication as well as resource requirements. The size directly affects the number of fingerprint records in the deduplication database as well as temporary space required for sorting these records. A smaller chunk size results in a large number of fingerprints and hence requires a significant amount of space for the deduplication database.

While the amount of storage that you save after deduplication depends heavily on the dataset and distribution of duplicates within the dataset, the chunk size can also affect the savings significantly. You must understand your dataset to get the best results after deduplication. A general rule of thumb is that a smaller chunk size saves more storage. A smaller chunk size results in more granular fingerprints and in general results in identifying more duplicates. However, smaller chunks have additional costs in terms of database size, deduplication time, and, more importantly, fragmentation. The deduplication database size can be significantly large for small chunk sizes. Higher fragmentation normally results in more file system metadata and hence can require more storage. The space consumed by the deduplication database and the increased file system metadata can reduce the savings achieved via deduplication. Additionally, fragmentation can also have a negative effect on performance. The Veritas File System (VxFS) deduplication algorithms try to reduce fragmentation by coalescing multiple contiguous duplicate chunks.

Larger chunk sizes normally result in a smaller deduplication database size, faster deduplication, and less fragmentation. These benefits sometimes come at the cost of less storage savings. If you have a large number duplicate files that are small in size, you still can choose a chunk size that is larger than the file size. A larger chunk size does not affect the deduplication of files that are smaller than the chunk size. In such cases, the fingerprint is calculated on the whole file, and the files are still deduplicated.

Symantec recommends a chunk size of 4k for Symantec VirtualStore, where multiple copies of virtual machine images are accessed over NFS. For all other datasets, Symantec recommends a chunk size of 16k or higher.

The space consumed by the deduplication database is a function of the amount of data in the file system and the deduplication chunk size. The space consumed by the deduplication database grows with time as new data is added to file system. Additional storage is required for temporary use, such as sorting fingerprints. The temporary storage may be freed after the work completes. Ensure that sufficient free space is available for deduplication to complete successfully. The deduplication might not start if the file system free space is less than approximately 15%. The deduplication sometimes needs more than 15% free space for smaller chunk sizes. In general, the space consumed reduces significantly with larger chunk sizes. Symantec recommends that you have approximately 20% free space for 4k chunks.

Deduplication and file system performance

Veritas File System (VxFS) deduplication uses shared extents to save storage when duplicate data is identified. Shared extents can significantly boost read performance for certain types of applications. These benefits are the result of the innovative use of file system page cache for data that resides in shared extents. Symantec VirtualStore, which serves a large number of virtual machine images, sees significant performance benefits from using shared extents.

The description of the FileSnaps feature contains more information about shared extents.

See [“About FileSnaps”](#) on page 580.

In general, any application or set of applications that read data residing in shared extents via multiple files are expected to have better read performance.

About the deduplication scheduler

The deduplication scheduler is a daemon that runs on all nodes and is responsible for deduplicating data as per the user-specified schedule. The scheduler is started on a node when you enable deduplication on the file system, but thereafter you must start the scheduler manually if you previously stopped the scheduler. Each file system can have its own schedule. The schedule and other configuration information for a given file system is stored within the file system. The location of the configuration file is `lost+found/dedup/local_config`.

The scheduler checks the configuration file every 30 minutes for changes and incorporates the changes if there are any. This periodic check also looks for newly mounted file systems. You can incorporate configuration changes immediately by restarting the scheduler.

When using the scheduler to deduplicate a file system's data automatically, the evaluation of changes in the file system is done by the File Change Log (FCL) feature. Scheduling deduplication to occur too infrequently in terms of days can cause the FCL to roll over, and thus the FCL feature can miss changes to the file system.

Symantec recommends that you schedule deduplication when the system activity is low. This ensures that the scheduler does not interfere with the regular workload of the system.

Deduplicating data

You deduplicate data using the `fsdedupadm` command. The `fsdedupadm` command performs the following functions:

Functionality	Command syntax
Enable the deduplication of a file system.	<code>fsdedupadm enable -c <i>chunk_size</i> [-q] <i>mount_point</i></code>
Disable the deduplication of a file system.	<code>fsdedupadm disable [-q] <i>mount_point</i></code>
Query the deduplication configuration of a file system.	<code>fsdedupadm list <i>mount_point</i> all</code>
Start a deduplication run on a file system.	<code>fsdedupadm start [-s] [-q] <i>mount_point</i></code>
Stop a deduplication run on a file system.	<code>fsdedupadm stop [-q] <i>mount_point</i></code>
Query the deduplication status of a file system.	<code>fsdedupadm status <i>mount_point</i> all</code>
Enable or disable the skipping of shared extents.	<code>fsdedupadm skipshared {true false} <i>mount_point</i></code>
Set the node on which the scheduled deduplication job will run.	<code>fsdedupadm setnodelist <i>nodelist</i> <i>mount_point</i> all</code>
Set the deduplication schedule for a file system.	<code>fsdedupadm setschedule <i>time</i> <i>mount_point</i></code>
Initiate a deduplication dry run on a file system.	<code>fsdedupadm dryrun [-o <i>threshold=#</i>] <i>mount_point</i></code>

Functionality**Command syntax**

Remove the deduplication configuration file and deduplication database on a file system.

`fsdedupadm remove mount_point`

For more information about the keywords, see the `fsdedupadm(1M)` manual page.

Enabling and disabling deduplication on a file system

You must enable deduplication on a file system by using the `fsdedupadm enable` command before you can use any of the deduplication functionality.

The following example enables deduplication on the file system mounted at `/mnt1`, and specifies a chunk size of 4096 for deduplication:

```
# /opt/VRTS/bin/fsdedupadm enable -c 4096 /mnt1
```

You can disable deduplication on a file system by using the `fsdedupadm disable` command.

The following example disables deduplication on the file system mounted at `/mnt1`:

```
# /opt/VRTS/bin/fsdedupadm disable /mnt1
```

Scheduling deduplication of a file system

You can set a schedule to deduplicate a file system automatically by using the `fsdedupadm setschedule` command. You can specify two categories of schedule options: run periodicity, and type periodicity. The granularity of schedule is limited to the time of day and the day of the month. The `fsdedupadm` command applies any relevant File Change Log tunables when setting the schedule.

See [“File Change Log administrative interface”](#) on page 931.

You must enable deduplication on the file system before you can set a schedule.

See [“Enabling and disabling deduplication on a file system”](#) on page 792.

You can schedule the deduplication run every hour or every specified number of hours, and every day or every specified number of days. You can also schedule the actual deduplication run to occur each time, or every specified number of times that the scheduled time elapses. During times that deduplication does not occur, the deduplication run only updates the fingerprints in the database.

The schedule commands are not cumulative. If a deduplication schedule comes up while the previous deduplication process is running for any reason, the upcoming deduplication is discarded and an warning message displays.

You can remove a schedule by specifying an empty string enclosed by double quotes (") for the schedule.

See the `fsdedupadm(1M)` manual page.

In the following example, deduplication for the file system `/vx/fs1` will be done at midnight, every other day:

```
# fsdedupadm setschedule "0 */2" /vx/fs1
```

In the following example, deduplication for the file system `/vx/fs1` will be done twice every day, once at midnight and once at noon:

```
# fsdedupadm setschedule "0,12 *" /vx/fs1
```

In the following example, deduplication for the file system `/vx/fs1` will be done four times every day, but only the fourth deduplication run will actually deduplicate the file system. The other runs will do the scanning and processing. This option achieves load distribution not only in a system, but also across the cluster.

```
# fsdedupadm setschedule "0,6,12,18 * 4" /vx/fs1
```

The following example removes the deduplication schedule from the file system `/vx/fs1`:

```
# fsdedupadm setschedule "" /vx/fs1
```

Performing a deduplication dry run

You can perform a dry run to determine the space savings of deduplication without actually modifying the file system. You must enable deduplication on the file system before you can perform a dry run.

See [“Enabling and disabling deduplication on a file system”](#) on page 792.

The following command initiates a deduplication dry run on the file system `/mnt1`:

```
# fsdedupadm dryrun /mnt1
```

You can specify `fsdedupadm` to perform the actual deduplication by specifying the `-o threshold` option. In this case, the `fsdedupadm` command performs an actual deduplication run if the expected space savings meets the specified threshold.

The following command initiates a deduplication dry run on the file system `/mnt1`, and performs the actual deduplication if the expected space savings crosses the threshold of 60:

```
# fsdedupadm dryrun -o threshold=60 /mnt1
```

Specifying the `-o threshold` option causes the `fsdedupadm` command to take Storage Checkpoints and enable the File Change Log for the file system.

Querying the deduplication status of a file system

You can query the deduplication status of a file system by using the `fsdedupadm status` command.

You must enable deduplication on the file system before you can query the deduplication status.

See [“Enabling and disabling deduplication on a file system”](#) on page 792.

The following command queries the deduplication status of the file system `/mnt1`:

```
# fsdedupadm status /mnt1
```

The following command queries the deduplication status of all running deduplication jobs:

```
# fsdedupadm status all
```

Starting and stopping the deduplication scheduler daemon

The state of the deduplication scheduler daemon, `fsdedupschd`, is maintained across reboots. If you started the `fsdedupschd` daemon prior to a reboot, the daemon is automatically restarted after the reboot. If you stopped the `fsdedupschd` daemon prior to a reboot, it remains stopped after the reboot. The default `fsdedupschd` daemon state is stopped.

You must enable deduplication on the file system before you can start or stop the scheduler daemon.

See [“Enabling and disabling deduplication on a file system”](#) on page 792.

The following command starts the `fsdedupschd` daemon:

```
# svcadm enable fsdedupschd
```

The following command stops the `fsdedupschd` daemon:

```
# svcadm disable fsdedupschd
```

Example of deduplicating a file system

The following example creates a file system, creates duplicate data on the file system, and deduplicates the file system.

To deduplicate a file system

- 1 Create the file system `fsvol1`:

```
# mkfs -F vxfs /dev/vx/rdisk/fsdg/fsvol1
```

- 2 Mount the file system as `/mnt1`:

```
# mount -F vxfs /dev/vx/dsk/fsdg/fsvol1 /mnt1
```

- 3 Make a temporary directory, `temp1`, on `/mnt1` and copy the `file1` file into the directory:

```
# mkdir /mnt1/temp1
# cd /mnt1/temp1
# cp /root/file1 .
```

```
# /opt/VRTS/bin/fsadm -S shared /mnt1
```

Mountpoint	Size(KB)	Available(KB)	Used(KB)	Logical_Size(KB)	Shared
/mnt1	4194304	3849468	283852	283852	0%

The `file1` file is approximately 250 MB, as shown by the output of the `fsadm` command.

- 4 Make another temporary directory, `temp2`, and copy the same file, `file1`, into the new directory:

```
# mkdir /mnt1/temp2
# cd /mnt1/temp2
# cp /root/file1 .
```

```
# /opt/VRTS/bin/fsadm -S shared /mnt1
```

Mountpoint	Size(KB)	Available(KB)	Used(KB)	Logical_Size(KB)	Shared
/mnt1	4194304	3588700	548740	548740	0%

By copying the same file into `temp2`, you now have duplicate data. The output of the `fsadm` command show that you are now using twice the amount of space.

5 Enable deduplication on the mount point /mnt1:

```
# /opt/VRTS/bin/fsdedupadm enable -c 4096 /mnt1
# /opt/VRTS/bin/fsdedupadm list /mnt1
```

Chunksize	Enabled	Schedule	NodeList	Priority	Filesystem
4096	YES	NONE	node1.company1.com	low	/mnt1

6 Start a deduplication run on the mount point /mnt1:

```
# /opt/VRTS/bin/fsdedupadm start /mnt1
UX:vxfs fsdedupadm: INFO: V-3-20: 0000: deduplication is started
on /mnt1.
```

7 Check status of deduplication:

```
# /opt/VRTS/bin/fsdedupadm status /mnt1
```

Saving	Status	Node	Type	Filesystem
74%	COMPLETED	node1.company1.com	MANUAL	/mnt1

2011/07/04 10:56:05 End detecting duplicates and filesystem changes.
Status 0

8 Verify that the file system was deduplicated by checking how much space you are using:

```
# /opt/VRTS/bin/fsadm -S shared /mnt1
```

Mountpoint	Size(KB)	Available(KB)	Used(KB)	Logical_Size(KB)	Shared
/mnt1	4194304	3834364	299136	566176	47%

The output shows that the used space is nearly identical to when you had only one copy of the `file1` file on the file system.

Deduplication results

The nature of the data is very important for deciding whether to enable deduplication. Databases or media files, such as JPEG, MP3, and MOV, might not be the best candidates for deduplication, as they have very little or no duplicate data. Virtual machine boot image files (`vmdk` files), user home directories, and file system with multiple copies of files are good candidates for deduplication. While smaller deduplication chunk size normally results into higher storage saving, it takes longer to deduplicate and requires a larger deduplication database.

Deduplication supportability

Veritas File System (VxFS) supports deduplication in the 6.0 release and later, and on file system disk layout version 9 and later. Deduplication is available on Linux (Red Hat and SuSE), Solaris SPARC, AIX and HP-UX Itanium.

Deduplication use cases

The following list includes several cases for which you would want to use the deduplication feature:

Home directories	User home directories often have multiple versions of the same files or file that have similar content, and therefore have redundant data that you can deduplicate.
Source code directories	Source code repositories usually have multiple files with incremental changes. The data that does not change from one file to the next can be deduplicated.
vmdk files	Once several virtual machines are cloned by using the FileSnap feature, the cloned virtual machines are subjected to operating system and security patches over their lifetime. As individual virtual machines cloned from a common source—the golden image—deviate from the source as a result of such activity, there is large amount of common content between them. Over time, this results in the loss of the initial storage savings. Deduplication of the new blocks added to these files restores the storage savings.

Deduplication limitations

The deduplication feature has the following limitations:

- A full backup of a deduplicated Veritas File System (VxFS) file system can require as much space in the target as a file system that has not been deduplicated. For example, if you have 2 TB of data that occupies 1 TB worth of disk space in the file system after deduplication, this data requires 2 TB of space on the target to back up the file system, assuming that the backup target does not do any deduplication. Similarly, when you restore such a file system, you must have 2 TB on the file system to restore the complete data. However, this freshly restored file system can be deduplicated again to regain the space savings. After a full file system restore, Symantec recommends that you remove any existing deduplication configuration using the `fsdedupadm remove` command and that you reconfigure deduplication using the `fsdedupadm enable` command.

- Deduplication is limited to a volume's primary filesset.
- Deduplication does not support mounted clone and snapshot mounted file system.
- After you restore data from a backup, you must deduplicate the restored data to regain any space savings provided by deduplication.
- If you use the cross-platform data sharing feature to convert data from one platform to another, you must remove the deduplication configuration file and database, re-enable deduplication, and restart deduplication after the conversion. The following example shows the commands that you must run, and you must run the commands in the order shown:

```
# /opt/VRTS/bin/fsdedupadm remove /mnt1
# /opt/VRTS/bin/fsdedupadm enable /mnt1
# /opt/VRTS/bin/fsdedupadm start /mnt1
```
- You cannot use the FlashBackup feature of NetBackup in conjunction with the data deduplication feature, because FlashBackup does not support disk layout Version 8 and 9.

Compressing files

This chapter includes the following topics:

- [About compressing files](#)
- [Compressing files with the vxcompress command](#)
- [Interaction of compressed files and other commands](#)
- [Interaction of compressed files and other features](#)
- [Interaction of compressed files and applications](#)
- [Use cases for compressing files](#)

About compressing files

Compressing files reduces the space used, while retaining the accessibility of the files and being transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files. Reads cause data to be uncompressed in memory, only; the on-disk copy of the file remains compressed. In contrast, after a write, the new data is uncompressed on disk.

Only user data is compressible. You cannot compress Veritas File System (VxFS) metadata.

After you compress a file, the inode number does not change, and file descriptors opened before the compression are still valid after the compression.

Compression is a property of a file. Thus, if you compress all files in a directory, for example, any files that you later copy into that directory do not automatically get compressed. You can compress the new files at any time by compressing the files in the directory again.

You compress files with the `vxcompress` command.

See [“Compressing files with the vxcompress command”](#) on page 801.

See the `vxcompress(1)` manual page.

To compress files, you must have VxFS file systems with disk layout Version 9 or later.

Note: When you back up compressed files to tape, VxFS stores the data in an uncompressed format. The files are uncompressed in memory and subsequently written to the tape. This results in increased CPU and memory usage when you back up compressed files.

About the compressed file format

A compressed file is a file with compressed extents. A `vxcompress` call compresses all extents of a file. However, writes to the file cause the affected extents to get uncompressed; the result can be files with both compressed and uncompressed extents.

About the file compression attributes

When you compress a file with the `vxcompress` command, `vxcompress` attaches the following information to the inode:

- Compression algorithm
- Compression strength, which is a number from 1 to 9
- Compression block size

This information is referred to as the file compression attributes. The purpose of the attributes are to collect the parameters used to create the compressed file. The information can then be read by a backup program.

The file compression attributes guarantee that a particular compressed file can only use one type and strength of compression. Recompressing a file using different attributes fails. To change the file compression attributes, you must explicitly uncompress first, and then recompress with the new options, even in the case where all extents are already uncompressed.

The file compression attributes do not indicate if all extents are compressed. Some extents might be incompressible, and other extents or even all extents might be uncompressed due to writes, but the file compression attributes remain. Only an explicit file uncompression can remove the attributes.

About the file compression block size

The file compression algorithm compresses data in the specified block size, which defaults to 1MB. Each compression block has its own extent descriptor in the inode. If the file or the last extent is smaller than the compression block size, then that smaller size gets compressed. The maximum block size is 1MB.

Extents with data that cannot be compressed are still marked as compressed extents. Even though such extents could not be compressed, marking these extents as compressed allows successive compression runs to skip these extents to save time. Shared extents cannot be compressed and do not get marked as compressed. Since the file compression algorithm looks at fixed-size blocks, the algorithm finds these incompressible extents in units of the file compression block size.

Compressing files with the vxcompress command

You can compress files with the `vxcompress` command. The `vxcompress` command performs the following functions:

Functionality	Command syntax
Compress files or directory trees	<code>vxcompress [-r] file_or_dir ...</code>
Uncompress files or directory trees	<code>vxcompress -u [-r] file_or_dir ...</code>
Report the compression savings in a file or directory tree	<code>vxcompress {-l -L} [-r] file_or_dir ...</code>
List the supported compression algorithms	<code>vxcompress -a</code>

See the `vxcompress(1)` manual page.

You can specify one or more filenames. If you specify the `-r` option, then you can specify directories, and the `vxcompress` command operates recursively on the directories.

You can specify the file compression algorithm and strength with the `vxcompress -t` command. The default algorithm is `gzip`, which is currently the only supported algorithm. The strength is a number from 1 to 9, with a default of 6. Strength 1 gives the fastest performance with least compression, while strength 9 gives the slowest performance with the greatest compression. For example, you specify strength 3 `gzip` compression as `"gzip-3"`.

When reporting the compression details for a file, the `vxcompress -l` command or `vxcompress -L` command displays the following information:

- Compression algorithm
- Strength
- Compression block size
- % of file data saved by compression
- % of extents that are compressed

This is the percentage of extents in the file that are compressed, without regard to the size of the extents. This percentage provides an idea of whether it is worthwhile to recompress the file. After recompression, the percentage is always 100%. However, shared extents are counted as uncompressed, and thus the percentage will be less than 100% if the file has shared extents.

If you attempt to compress a file with the `vxcompress` command and the extents have data that cannot be compressed, the command still marks the file as compressed and replaces the extents with compressed extent descriptors.

If you recompress a file, you do not need to specify any options with the `vxcompress` command. The command automatically uses the options that you used to compress the file previously.

Examples of using the `vxcompress` command

The following command compresses the `file1` file, using the default algorithm and strength of gzip-6:

```
$ vxcompress file1
```

The following command recursively compresses all files below the `dir1` directory, using the gzip algorithm at the highest strength (9):

```
$ vxcompress -r -t gzip-9 dir1
```

The following command compresses the `file2` file and all files below the `dir2` directory, using the gzip algorithm at strength 3, while limiting the `vxcompress` command to a single thread and reducing the scheduling priority:

```
$ vxcompress -r -t gzip-3 file2 dir1
```

The following command displays the results of compressing the `file1` file in human-friendly units:

```
$ vxcompress -L file1
%Comp   Physical   Logical   %Exts   Alg-Str   BSize   Filename
  99%         1 KB      159 KB   100%    gzip-6   1024k   file1
```

The following command uncompresses the `file1` file:

```
$ vxcompress -u file1
```

Interaction of compressed files and other commands

Table 36-1 describes how compressed files interact with other Veritas Storage Foundation commands and base operating system commands.

Table 36-1

Command	Interaction with compressed files
<code>df</code>	The <code>df</code> command shows the actual blocks in use by the file system. This number includes the compression savings, but the command does not display the savings explicitly. See the <code>df(1)</code> manual page.
<code>du</code>	The <code>du</code> usually uses the block count and thus implicitly shows the results of compression, but the GNU <code>du</code> command has an option to use the file size instead, which is not changed by compression. See the <code>du(1)</code> manual page.
<code>fsadm -S</code>	The <code>fsadm -S compressed</code> command can report the space savings due to compressed files. See the <code>fsadm_vxfs(1)</code> manual page.
<code>fsmmap -p</code>	The <code>fsmmap</code> command can report information on compressed and uncompressed extents with the <code>-p</code> option. The reported logical size is the size of the uncompressed data, while the reported extent size is the size of the compressed data on disk. For compressed extents, the two sizes might differ. See the <code>fsmmap(1)</code> manual page.
<code>ls -l</code> <code>ls -s</code>	The inode size reported by a <code>stat</code> call is the logical size, as shown by the <code>ls -l</code> command. This size is not affected by compression. On the other hand, the block count reflects the actual blocks used. As such, the <code>ls -s</code> command shows the result of compression. See the <code>ls(1)</code> manual page.

Table 36-1 (continued)

Command	Interaction with compressed files
<code>vxdump</code>	The <code>vxdump</code> command uncompresses compressed extents as it encounters them, meaning that compression is not preserved across a backup or restore operation.
<code>vxquota</code>	Your quota usage decreases based on the space saved due to compression. See the <code>vxquota(1M)</code> manual page.

Interaction of compressed files and other features

Table 36-2 describes how compressed files interact with other Veritas Storage Foundation features.

Table 36-2

Feature	Interaction with compressed files
Cross-Platform Data Sharing	If you convert a disk or file system from one platform that supports compression to a platform that does not support compression and the file system contains compressed files, the <code>fsccdsconv</code> command displays a message that some files violate the CDS limits and prompts you to confirm if you want to continue the conversion. If you continue, the conversion completes successfully, but the compressed files will not be accessible on the new platform.
File Change Log	The File Change Log feature does not detect file compressions nor uncompressions.
Shared extents (FileSnap and, on Solaris SPARC, deduplication)	Shared extents do not get compressed.
SmartTier	The SmartTier feature does not support compression. A placement policy cannot move existing compressed extents. Newly-allocated compressed extents follow the existing placement policy.
Space reservation (<code>setext -r</code>)	When a file is compressed, any space reserved via the <code>setext -r</code> command beyond the end-of-file is discarded, and is not restored when the file is uncompressed. The <code>setext -r</code> command cannot be used to reserve space for files that are compressed.

Table 36-2 (continued)

Feature	Interaction with compressed files
Storage Checkpoints	If a file system contains compressed files and you create a Storage Checkpoint of that file system, you can access those files normally through the Storage Checkpoint. However, you cannot compress nor uncompress a file that is already in a mounted Storage Checkpoint.

Interaction of compressed files and applications

In general, applications notice no difference between compressed and uncompressed files, although reads and writes to compressed extents are slower than reads and writes to uncompressed extents. When an application reads a compressed file, the file system does not perform its usual readahead to avoid the CPU load that this can require. However, when reading from the primary fileset, the file system uncompresses an entire compression block (default 1 MB) and leaves these pages in the page cache. Thus, sequential reads of the file usually only incur an extra cost when crossing a compression block boundary. The situation is different when reading from a file in a Storage Checkpoint; in this case, nothing goes into the page cache beyond the data actually requested. For optimal read performance of a compressed file accessed through a Storage Checkpoint, the application should use a read size that matches the compression block size.

When writing to compressed extents, ensure that you have sufficient disk space and disk quota limits for the new uncompressed extents since the write uncompresses the extents. If you do not have sufficient disk space, the write can fail with the ENOSPC error. If you do not have enough disk quota, the write can fail with the EDQUOT error.

An application that reads data from a compressed file and then copies the file elsewhere, such as `tar`, `cpio`, `cp`, or `vi`, does not preserve compression in the new data. The same is true of some backup programs.

Backup programs that read file data through the name space do not notice that the file is compressed. The backup program receives uncompressed data, and the compression is lost.

You cannot use the FlashBackup feature of NetBackup in conjunction with the file compression feature, because FlashBackup does not support disk layout Version 8 and 9.

Use cases for compressing files

The following list contains common use case categories:

- [Compressed files and databases](#)
- [Compressing all files that meet the specified criteria](#)

Compressed files and databases

Compressing files helps to reduce the storage cost in a database environment. For Oracle databases, compression provides an excellent value add to reduce storage cost for archived logs, partitioned tables, and infrequently accessed tablespaces and datafiles. The compression ratio of database files depends on the type of object stored in the datafiles. Oracle traditionally stores TABLES and INDEXES in datafiles, in which case the compression ratio depends on type of columns associated with the TABLE and the type of keys in the INDEXES. Oracle also has the ability to store unstructured data, such as XML, spreadsheets, MS Word documents, and pictures, within a TABLE via the Secured Files feature. These types of unstructured data are very good candidates for compression. You can achieve up to 90% compression for archived logs, and about 50% to 65% compression for Oracle datafiles and indexes.

Oracle database files can be compressed and uncompressed as needed while the database is active, although this can have a significant performance impact on the database. Other than reduced I/O response time, compression runs seamlessly while the Oracle database is online and actively doing transactions to the files. Compression works seamlessly with advanced I/O methods, such as direct I/O, asynchronous I/O, concurrent I/O, ODM, and Cached ODM. Any updates and new inserts to the datafile result in uncompressing the portion of the file associated with the write. The queries get uncompressed data in memory and the file remains compressed.

Note: You cannot compress Quick I/O files.

You can run the `vxcompress` command as a DBA user.

The following use cases apply to databases:

- [Supported database versions and environment](#)
- [Compressing archive logs](#)
- [Compressing read-only tablespaces](#)
- [Compressing infrequently accessed table partitions](#)

- [Compressing infrequently accessed datafiles](#)
- [Best practices for compressing files in an Oracle database](#)

Supported database versions and environment

You can use compressed files with Oracle versions 10gR2, 11gR1, and 11gR2. Compression is supported in Veritas Storage Foundation for Oracle RAC (SFRAC) and Veritas Storage Foundation Cluster File System High Availability (SFCFSHA). In a clustered environment, such as SFRAC and SFCFSHA, Symantec recommends that you compress files on a node that has minimal load. In a Fast Failover SFCFSHA environment, Symantec recommends that you compress files on a passive node where the database is offline.

Compressing archive logs

Archive logs are critical files required for database recovery. In a busy online transaction processing (OLTP) database, several gigabytes of archive logs are generated each day. Company guidelines often mandate preserving archive logs for several days. The Oracle archive logs are read-only files and are never updated after they are generated. During recovery, Oracle reads archive logs sequentially. As such, archive logs are very good candidates for compression, and archive logs are highly compressible.

The following example procedure compresses all archive logs that are older than a day.

To compress all archive logs that are older than a day

- 1 As an Oracle DBA, run the following query and get the archive log location:

```
$ SQL> select destination from v$archive_dest where status = 'VALID'
and valid_now = 'YES';
```

Assume `/oraarch/MYDB` is the archive log destination.

- 2 Compress all of the archive logs that are older than a day:

```
$ find /oraarch/MYDB -mtime +1 -exec /opt/VRTS/bin/vxcompress {} \;
```

You can run this step daily via a scheduler, such as `cron`.

Compressing read-only tablespaces

In a large database environment, it is a common practice to keep static tablespaces that do not have any changes in read-only mode. The primary purpose of read-only tablespaces is to eliminate the need to perform backup and recovery of large, static portions of a database. Read-only tablespaces also provide a way to protecting

historical data so that users cannot modify it. Making a tablespace read-only prevents updates on all tables and objects residing in the tablespace, regardless of a user's update privilege level. These kinds of read-only tablespaces are excellent candidates for compression. In some cases such as month end reports, there may be large queries executed against these read-only tablespaces. To make the report run faster, you can uncompress the tablespace on demand before running the monthly reports.

In the following example, a sporting goods company has its inventory divided into two tablespaces: `winter_items` and `summer_items`. In the end of the Spring season, you can compress the `winter_item` tablespace and uncompress the `summer_item` tablespace. You can do the reverse actions at end of the Summer season. The following example procedure performs these tasks.

To compress and uncompress tablespaces depending on the season

- 1 Using SQL, get a list of files in each tablespace and store the result in the files `summer_files` and `winter_files`.

```
$ SQL> select file_name from dba_data_files where  
tablespace_name = 'WINTER_ITEM';  
Store the result in winter_files  
$ SQL> select file_name from dba_data_files where  
tablespace_name = 'SUMMER_ITEM';  
Store the result in summer_files
```

- 2 Compress the `winter_files` file:

```
$ /opt/VRTS/bin/vxcompress `/bin/cat winter_files`
```

- 3 Uncompress the `summer_files` file:

```
$ /opt/VRTS/bin/vxcompress -u `/bin/cat summer_files`
```

Compressing infrequently accessed table partitions

Partitioned tables is a frequently used feature for large Oracle databases. Table partitioning improves database queries and updates because partitioning helps parallelizing transactions that use Parallel Queries. Partitioning also makes maintenance of database easy and improves the availability of TABLES. If a partition is down, only the corresponding portion of the TABLE goes offline and the rest of the TABLE remains online. In a telecommunications environment, a common practice is to partition a 'call_details' table by month or quarter. The contents in the partition are less active as the partition gets older. The new records are added to a new partition, and previous quarter records do not get updated. Since telecommunications

databases are generally very large, compressing last year's data provides great savings.

In the following example, assume that the table 'CALL_DETAIL' is partitioned by quarters, and the partition names are CALL_2010_Q1, CALL_2010_Q2, and CALL_2011_Q1, and so on. In the first Quarter of 2011, you can compress the CALL_2010_Q1 data.

To compress the CALL_2010_Q1 partition

- 1 Use SQL to retrieve the filenames belonging to the CALL_2010_Q1 partition:

```
$ SQL> select tablespace_name from dba_tab_partitions  
where table_name = 'CALL_DETAIL' and partition_name = 'CALL_2010_Q1';
```

Assume that the query returns "TBS_2010_Q1".

- 2 Store the names in the my_compress_files file:

```
$ SQL> select file_name from dba_data_files where  
tablespace_name = 'TBS_2010_Q1';  
Store the result in my_compress_files.
```

- 3 Compress the files:

```
$ /opt/VRTS/bin/vxcompress `/bin/cat my_compress_files`
```

Compressing infrequently accessed datafiles

Many customer databases do not use the Oracle partitioning feature. If partitioning is not used, then you can use Oracle catalog queries to identify datafiles that are not very active. Periodically, you can query the catalog tables and identify the least active datafiles and compress those files, as illustrated in the following example procedure.

To identify the least active datafiles and compress those files

- 1 Query v\$filestat and identify the least active datafiles:

```
$ SQL> select name, phylds + phywrts 'TOT_IO' from v$datafile d
and v$filestat f where d.file# = f.file# order by TOT_IO;
```

- 2 Select files that have the least I/O activity from the report and compress those files:

```
$ /opt/VRTS/bin/vxcompress file1 file2 file3 ...
```

- 3 Periodically run the query again to ensure that the compressed files do not have increased I/O activity. If I/O activity increases, uncompress the files:

```
$ /opt/VRTS/bin/vxcompress -u file1 file2 file3 ...
```

Best practices for compressing files in an Oracle database

Even though an Oracle database runs without any errors when files are compressed, increased I/O to compressed files decreases database performance. Use the following guidelines for compressing Oracle datafiles:

- Do not compress database control files.
- Do not compress files belonging to TEMPORARY tablespaces.
- Do not compress files belonging to SYSTEM and SYSAUX tablespace.
- Monitor the I/O activity on compressed files periodically and uncompress the files if I/O activity increases.

Compressing all files that meet the specified criteria

You can find all files that meet the specified criteria and pipe the results to the `vxcompress` command to compress all of those files. The following example compresses all files in `/mnt` that have not been modified for more than 30 days:

```
$ find /mnt -mtime +30 | xargs /opt/VRTS/bin/vxcompress
```

Reference

- [Chapter 37. Managing storage](#)
- [Chapter 38. Rootability](#)
- [Chapter 39. Managing tunable parameters](#)
- [Chapter 40. Quotas](#)
- [Chapter 41. File Change Log](#)
- [Chapter 42. Reverse path name lookup](#)
- [Appendix A. Disk layout](#)
- [Appendix B. Command reference](#)
- [Appendix C. Creating a starter database](#)

Managing storage

This chapter includes the following topics:

- [Moving volumes or disks](#)
- [Monitoring and controlling tasks](#)
- [Using vxnotify to monitor configuration changes](#)
- [Performing online relayout](#)
- [Adding a mirror to a volume](#)
- [Configuring SmartMove](#)
- [Removing a mirror](#)
- [Setting tags on volumes](#)
- [Managing disk groups](#)
- [Managing volumes](#)
- [Managing plexes and subdisks](#)
- [Decommissioning storage](#)

Moving volumes or disks

This section describes moving volumes or disks.

Moving volumes from a VM disk

Before you disable or remove a disk, you can move the data from that disk to other disks on the system that have sufficient space.

To move volumes from a disk

- 1 From the `vxdiskadm` main menu, select `Move volumes from a disk`.
- 2 At the following prompt, enter the disk name of the disk whose volumes you want to move, as follows:

```
Enter disk name [<disk>,list,q,?] mydg01
```

You can now optionally specify a list of disks to which the volume(s) should be moved. At the prompt, do one of the following:

- Press **Enter** to move the volumes onto available space in the disk group.
- Specify the disks in the disk group that should be used, as follows:

```
Enter disks [<disk ...>,list]
```

```
VxVM NOTICE V-5-2-283 Requested operation is to move all  
volumes from disk mydg01 in group mydg.
```

NOTE: This operation can take a long time to complete.

```
Continue with operation? [y,n,q,?] (default: y)
```

As the volumes are moved from the disk, the `vxdiskadm` program displays the status of the operation:

```
VxVM vxevac INFO V-5-2-24 Move volume voltest ...
```

When the volumes have all been moved, the `vxdiskadm` program displays the following success message:

```
VxVM INFO V-5-2-188 Evacuation of disk mydg02 is complete.
```

- 3 At the following prompt, indicate whether you want to move volumes from another disk (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Move volumes from another disk? [y,n,q,?] (default: n)
```

Moving disks between disk groups

To move a disk between disk groups, remove the disk from one disk group and add it to the other. For example, to move the physical disk `c0t3d0` (attached with the disk name `salesdg04`) from disk group `salesdg` and add it to disk group `mktdg`, use the following commands:

```
# vxdg -g salesdg rmdisk salesdg04  
# vxdg -g mkt dg adddisk mkt dg02=c0t3d0
```

Warning: This procedure does not save the configurations nor data on the disks.

You can also move a disk by using the `vxdiskadm` command. Select **Remove a disk** from the main menu, and then select **Add or initialize a disk**.

To move disks and preserve the data on these disks, along with VxVM objects, such as volumes:

See [“Moving objects between disk groups”](#) on page 821.

Reorganizing the contents of disk groups

There are several circumstances under which you might want to reorganize the contents of your existing disk groups:


- To group volumes or disks differently as the needs of your organization change. For example, you might want to split disk groups to match the boundaries of separate departments, or to join disk groups when departments are merged.
- To isolate volumes or disks from a disk group, and process them independently on the same host or on a different host. This allows you to implement off-host processing solutions for the purposes of backup or decision support.
- To reduce the size of a disk group’s configuration database in the event that its private region is nearly full. This is a much simpler solution than the alternative of trying to grow the private region.
- To perform online maintenance and upgrading of fault-tolerant systems that can be split into separate hosts for this purpose, and then rejoined.

Use the `vxdg` command to reorganize your disk groups.

The `vxdg` command provides the following operations for reorganizing disk groups:

- The `move` operation moves a self-contained set of VxVM objects between imported disk groups. This operation fails if it would remove all the disks from the source disk group. Volume states are preserved across the move.


[Figure 37-1](#) shows the `move` operation.

Figure 37-1 Disk group move operation

- The `split` operation removes a self-contained set of VxVM objects from an imported disk group, and moves them to a newly created target disk group. This operation fails if it would remove all the disks from the source disk group, or if an imported disk group exists with the same name as the target disk group. An existing deported disk group is destroyed if it has the same name as the target disk group (as is the case for the `vxvg init` command).


Figure 37-2 shows the `split` operation.

Figure 37-2 Disk group split operation



- The `join` operation removes all VxVM objects from an imported disk group and moves them to an imported target disk group. The source disk group is removed when the join is complete.

[Figure 37-3](#) shows the `join` operation.

Figure 37-3 Disk group join operation

These operations are performed on VxVM objects such as disks or top-level volumes, and include all component objects such as sub-volumes, plexes and subdisks. The objects to be moved must be self-contained, meaning that the disks that are moved must not contain any other objects that are not intended for the move.

For site-consistent disk groups, any of the move operations (move, split, and join) fail if the VxVM objects that are moved would not meet the site consistency conditions after the move. For example, a volume that is being moved may not have a plex on one of the sites configured in the target disk group. The volume would not meet the conditions for the `allsites` flag in the target disk group. Use the `-f` (force) option to enable the operation to succeed, by turning off the `allsites` flag on the object.

If you specify one or more disks to be moved, all VxVM objects on the disks are moved. You can use the `-o expand` option to ensure that `vxvg` moves all disks on which the specified objects are configured. Take care when doing this as the result may not always be what you expect. You can use the `listmove` operation with `vxvg` to help you establish what is the self-contained set of objects that corresponds to a specified set of objects.

Warning: Before moving volumes between disk groups, stop all applications that are accessing the volumes, and unmount all file systems that are configured on these volumes.

If the system crashes or a hardware subsystem fails, VxVM attempts to complete or reverse an incomplete disk group reconfiguration when the system is restarted or the hardware subsystem is repaired, depending on how far the reconfiguration had progressed. If one of the disk groups is no longer available because it has been imported by another host or because it no longer exists, you must recover the disk group manually.

See the *Veritas Storage Foundation and High Availability Troubleshooting Guide*.

Limitations of disk group split and join

The disk group split and join feature has the following limitations:

- Disk groups involved in a move, split or join must be version 90 or greater. See [“Upgrading the disk group version”](#) on page 848.
- The reconfiguration must involve an integral number of physical disks.
- Objects to be moved must not contain open volumes.
- Disks cannot be moved between CDS and non-CDS compatible disk groups.
- By default, VxVM automatically recovers and starts the volumes following a disk group move, split or join. If you have turned off the automatic recovery feature, volumes are disabled after a move, split, or join. Use the `vxrecover -m` and `vxvol startall` commands to recover and restart the volumes.
- Data change objects (DCOs) and snap objects that have been dissociated by Persistent FastResync cannot be moved between disk groups.
- Veritas Volume Replicator (VVR) objects cannot be moved between disk groups.
- For a disk group move to succeed, the source disk group must contain at least one disk that can store copies of the configuration database after the move.
- For a disk group split to succeed, both the source and target disk groups must contain at least one disk that can store copies of the configuration database after the split.
- For a disk group move or join to succeed, the configuration database in the target disk group must be able to accommodate information about all the objects in the enlarged disk group.
- Splitting or moving a volume into a different disk group changes the volume's record ID.

- The operation can only be performed on the master node of a cluster if either the source disk group or the target disk group is shared.
- In a cluster environment, disk groups involved in a move or join must both be private or must both be shared.
- If a cache object or volume set that is to be split or moved uses ISP volumes, the storage pool that contains these volumes must also be specified.

Listing objects potentially affected by a move

To display the VxVM objects that would be moved for a specified list of objects, use the following command:

```
# vxdg [-o expand] listmove sourcedg targetdg object ...
```

The following example lists the objects that would be affected by moving volume `vol1` from disk group `mydg` to `newdg`:

```
# vxdg listmove mydg newdg vol1
mydg01 c0t1d0s2 mydg05 clt96d0s2 vol1 vol1-01 vol1-02 mydg01-01
mydg05-01
```

However, the following command produces an error because only a part of the volume `vol1` is configured on the disk `mydg01`:

```
# vxdg listmove mydg newdg mydg01
VxVM vxdg ERROR V-5-2-4597 vxdg listmove mydg newdg failed
VxVM vxdg ERROR V-5-2-3091 mydg05 : Disk not moving, but
subdisks on it are
```

Specifying the `-o expand` option, as shown below, ensures that the list of objects to be moved includes the other disks (in this case, `mydg05`) that are configured in `vol1`:

```
# vxdg -o expand listmove mydg newdg mydg01
mydg01 c0t1d0s2 mydg05 clt96d0s2 vol1 vol1-01 vol1-02 mydg01-01 mydg05-01
```

Moving DCO volumes between disk groups


When you move the parent volume (such as a snapshot volume) to a different disk group, its DCO volume must accompany it. If you use the `vxassist addlog`, `vxmake` or `vxdco` commands to set up a DCO for a volume, you must ensure that the disks that contain the plexes of the DCO volume accompany their parent volume during the move. You can use the `vxprint` command on a volume to examine the configuration of its associated DCO volume.

If you use the `vxassist` command to create both a volume and its DCO, or the `vxsnap prepare` command to add a DCO to a volume, the DCO plexes are automatically placed on different disks from the data plexes of the parent volume. In previous releases, version 0 DCO plexes were placed on the same disks as the data plexes for convenience when performing disk group split and move operations. As version 20 DCOs support dirty region logging (DRL) in addition to Persistent FastResync, it is preferable for the DCO plexes to be separated from the data plexes. This improves the performance of I/O from/to the volume, and provides resilience for the DRL logs.


[Figure 37-4](#) shows some instances in which it is not be possible to split a disk group because of the location of the DCO plexes on the disks of the disk group.

See [“Volume snapshots”](#) on page 130.


Figure 37-4 Examples of disk groups that can and cannot be split




The disk group can be split as the DCO plexes are on dedicated disks, and can therefore accompany the disks that contain the volume data



The disk group cannot be split as the DCO plexes cannot accompany their volumes. One solution is to relocate the DCO plexes. In this example, use an additional disk in the disk group as an intermediary to swap the misplaced DCO plexes. Alternatively, to improve DRL performance and resilience, allocate the DCO plexes to dedicated disks.



The disk group can be split as the DCO plexes can accompany their volumes. However, you may not wish the data in the portions of the disks marked "?" to be moved as well.



The disk group cannot be split as this would separate the disks containing Volume 2's data plexes. Possible solutions are to relocate the snapshot DCO plex to the snapshot plex disk, or to another suitable disk that can be moved.

Moving objects between disk groups

To move a self-contained set of VxVM objects from an imported source disk group to an imported target disk group, use the following command:

```
# vxdg [-o expand] [-o override|verify] move sourcedg targetdg \
    object ...
```

The `-o expand` option ensures that the objects that are actually moved include all other disks containing subdisks that are associated with the specified objects or with objects that they contain.

The default behavior of `vxdg` when moving licensed disks in an EMC array is to perform an EMC disk compatibility check for each disk involved in the move. If the compatibility checks succeed, the move takes place. `vxdg` then checks again to ensure that the configuration has not changed since it performed the compatibility check. If the configuration has changed, `vxdg` attempts to perform the entire move again.

Note: You should only use the `-o override` and `-o verify` options if you are using an EMC array with a valid timefinder license. If you specify one of these options and do not meet the array and license requirements, a warning message is displayed and the operation is ignored.

The `-o override` option enables the move to take place without any EMC checking.

The `-o verify` option returns the access names of the disks that would be moved but does not perform the move.

The following output from `vxprint` shows the contents of disk groups `rootdg` and `mydg`.

The output includes two utility fields, `TUTILO` and `PUTILO`. VxVM creates these fields to manage objects and communications between different commands and Symantec products. The `TUTILO` values are temporary; they are not maintained on reboot.

The `PUTILO` values are persistent; they are maintained on reboot.

```
# vxprint
Disk group: rootdg
TY NAME      ASSOC      KSTATE    LENGTH    PLOFFS    STATE    TUTILO    PUTILO
dg rootdg    rootdg     -         -         -         -         -         -
dm rootdg02  c1t97d0s2 -         17678493 -         -         -         -
dm rootdg03  c1t112d0s2 -         17678493 -         -         -         -
dm rootdg04  c1t114d0s2 -         17678493 -         -         -         -
dm rootdg06  c1t98d0s2 -         17678493 -         -         -         -

Disk group: mydg
TY NAME      ASSOC      KSTATE    LENGTH    PLOFFS    STATE    TUTILO    PUTILO
dg mydg      mydg       -         -         -         -         -         -
dm mydg01    c0t1d0s2  -         17678493 -         -         -         -
```

dm mydg05	c1t96d0s2	-	17678493	-	-	-	-
dm mydg07	c1t99d0s2	-	17678493	-	-	-	-
dm mydg08	c1t100d0s2	-	17678493	-	-	-	-
v vol1	fsgen	ENABLED	2048	-	ACTIVE	-	-
pl vol1-01	vol1	ENABLED	3591	-	ACTIVE	-	-
sd mydg01-01	vol1-01	ENABLED	3591	0	-	-	-
pl vol1-02	vol1	ENABLED	3591	-	ACTIVE	-	-
sd mydg05-01	vol1-02	ENABLED	3591	0	-	-	-

The following command moves the self-contained set of objects implied by specifying `disk mydg01` from disk group `mydg` to `rootdg`:

```
# vxdg -o expand move mydg rootdg mydg01
```

By default, VxVM automatically recovers and starts the volumes following a disk group move. If you have turned off the automatic recovery feature, volumes are disabled after a move. Use the following commands to recover and restart the volumes in the target disk group:

```
# vxrecover -g targetdg -m [volume ...]
# vxvol -g targetdg startall
```

The output from `vxprint` after the move shows that not only `mydg01` but also volume `vol1` and `mydg05` have moved to `rootdg`, leaving only `mydg07` and `mydg08` in disk group `mydg`:

```
# vxprint
Disk group: rootdg
TY NAME      ASSOC      KSTATE    LENGTH    PLOFFS    STATE     TUTILO    PUTILO
dg rootdg    rootdg     -          -          -          -          -
dm mydg01    c0t1d0s2   -          17678493   -          -          -
dm rootdg02  c1t97d0s2   -          17678493   -          -          -
dm rootdg03  c1t112d0s2 -          17678493   -          -          -
dm rootdg04  c1t114d0s2 -          17678493   -          -          -
dm mydg05    c1t96d0s2   -          17678493   -          -          -
dm rootdg06  c1t98d0s2   -          17678493   -          -          -
v vol1       fsgen      ENABLED    2048       -          ACTIVE    -
pl vol1-01   vol1       ENABLED    3591       -          ACTIVE    -
sd mydg01-01 vol1-01    ENABLED    3591       0          -          -
pl vol1-02   vol1       ENABLED    3591       -          ACTIVE    -
sd mydg05-01 vol1-02    ENABLED    3591       0          -          -

Disk group: mydg
TY NAME      ASSOC      KSTATE    LENGTH    PLOFFS    STATE     TUTILO    PUTILO
dg mydg      mydg       -          -          -          -          -
```

```
dm mydg07      c1t99d0s2  -          17678493  -          -          -          -
dm mydg08      c1t100d0s2 -          17678493  -          -          -          -
```

The following commands would also achieve the same result:

```
# vxdg move mydg rootdg mydg01 mydg05
# vxdg move mydg rootdg voll
```

See [“Moving objects between shared disk groups”](#) on page 383.

Splitting disk groups

To remove a self-contained set of VxVM objects from an imported source disk group to a new target disk group, use the following command:

```
# vxdg [-o expand] [-o override|verify] split sourcedg targetdg \
object ...
```

See [“Moving objects between disk groups”](#) on page 821.

The following output from `vxprint` shows the contents of disk group `rootdg`.

The output includes two utility fields, `TUTILO` and `PUTILO`. VxVM creates these fields to manage objects and communications between different commands and Symantec products. The `TUTILO` values are temporary; they are not maintained on reboot. The `PUTILO` values are persistent; they are maintained on reboot.

```
# vxprint
Disk group: rootdg
TY NAME      ASSOC      KSTATE      LENGTH      PLOFFS      STATE      TUTILO      PUTILO
dg rootdg    rootdg     -           -           -           -           -           -
dm rootdg01  c0t1d0s2  -           17678493    -           -           -           -
dm rootdg02  c1t97d0s2 -           17678493    -           -           -           -
dm rootdg03  c1t112d0s2 -          17678493    -           -           -           -
dm rootdg04  c1t114d0s2 -          17678493    -           -           -           -
dm rootdg05  c1t96d0s2  -           17678493    -           -           -           -
dm rootdg06  c1t98d0s2  -           17678493    -           -           -           -
dm rootdg07  c1t99d0s2  -           17678493    -           -           -           -
dm rootdg08  c1t100d0s2 -          17678493    -           -           -           -
v voll      fsgen      ENABLED     2048        -           ACTIVE      -           -
pl voll-01  voll       ENABLED     3591        -           ACTIVE      -           -
sd rootdg01-01 voll-01     ENABLED     3591        0           -           -           -
pl voll-02  voll       ENABLED     3591        -           ACTIVE      -           -
sd rootdg05-01 voll-02     ENABLED     3591        0           -           -           -
```

The following command removes disks `rootdg07` and `rootdg08` from `rootdg` to form a new disk group, `mydg`:

```
# vxdg -o expand split rootdg mydg rootdg07 rootdg08
```

By default, VxVM automatically recovers and starts the volumes following a disk group split. If you have turned off the automatic recovery feature, volumes are disabled after a split. Use the following commands to recover and restart the volumes in the target disk group:

```
# vxrecover -g targetdg -m [volume ...]
# vxvol -g targetdg startall
```

The output from `vxprint` after the split shows the new disk group, `mydg`:

```
# vxprint
```

```
Disk group: rootdg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTIL0	PUTIL0
dg	rootdg	rootdg	-	-	-	-	-	-
dm	rootdg01	c0t1d0s2	-	17678493	-	-	-	-
dm	rootdg02	c1t97d0s2	-	17678493	-	-	-	-
dm	rootdg03	c1t112d0s2-	-	17678493	-	-	-	-
dm	rootdg04	c1t114d0s2-	-	17678493	-	-	-	-
dm	rootdg05	c1t96d0s2	-	17678493	-	-	-	-
dm	rootdg06	c1t98d0s2	-	17678493	-	-	-	-
v	vol1	fsgen	ENABLED	2048	-	ACTIVE	-	-
pl	vol1-01	vol1	ENABLED	3591	-	ACTIVE	-	-
sd	rootdg01-01	vol1-01	ENABLED	3591	0	-	-	-
pl	vol1-02	vol1	ENABLED	3591	-	ACTIVE	-	-
sd	rootdg05-01	vol1-02	ENABLED	3591	0	-	-	-

```
Disk group: mydg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTIL0	PUTIL0
dg	mydg	mydg	-	-	-	-	-	-
dm	rootdg07	c1t99d0s2	-	17678493	-	-	-	-
dm	rootdg08	c1t100d0s2-	-	17678493	-	-	-	-

See [“Splitting shared disk groups”](#) on page 383.

Joining disk groups

To remove all VxVM objects from an imported source disk group to an imported target disk group, use the following command:

```
# vxdg [-o override|verify] join sourcedg targetdg
```

See [“Moving objects between disk groups”](#) on page 821.

Note: You cannot specify `rootdg` as the source disk group for a `join` operation.

The following output from `vxprint` shows the contents of the disk groups `rootdg` and `mydg`.

The output includes two utility fields, `TUTILO` and `PUTILO`.. VxVM creates these fields to manage objects and communications between different commands and Symantec products. The `TUTILO` values are temporary; they are not maintained on reboot. The `PUTILO` values are persistent; they are maintained on reboot.

vxprint

Disk group: `rootdg`

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	rootdg	rootdg	-	-	-	-	-	-
dm	rootdg01	c0t1d0s2	-	17678493	-	-	-	-
dm	rootdg02	c1t97d0s2	-	17678493	-	-	-	-
dm	rootdg03	c1t112d0s2	-	17678493	-	-	-	-
dm	rootdg04	c1t114d0s2	-	17678493	-	-	-	-
dm	rootdg07	c1t99d0s2	-	17678493	-	-	-	-
dm	rootdg08	c1t100d0s2	-	17678493	-	-	-	-

Disk group: `mydg`

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	mydg	mydg	-	-	-	-	-	-
dm	mydg05	c1t96d0s2	-	17678493	-	-	-	-
dm	mydg06	c1t98d0s2	-	17678493	-	-	-	-
v	vol1	fsgen	ENABLED	2048	-	ACTIVE	-	-
pl	vol1-01	vol1	ENABLED	3591	-	ACTIVE	-	-
sd	mydg01-01	vol1-01	ENABLED	3591	0	-	-	-
pl	vol1-02	vol1	ENABLED	3591	-	ACTIVE	-	-
sd	mydg05-01	vol1-02	ENABLED	3591	0	-	-	-

The following command joins disk group `mydg` to `rootdg`:

```
# vxvg join mydg rootdg
```

By default, VxVM automatically recovers and starts the volumes following a disk group join. If you have turned off the automatic recovery feature, volumes are disabled after a join. Use the following commands to recover and restart the volumes in the target disk group:

```
# vxrecover -g targetdg -m [volume ...]
# vxvol -g targetdg startall
```

The output from `vxprint` after the join shows that disk group `mydg` has been removed:

```
# vxprint
Disk group: rootdg
TY NAME      ASSOC      KSTATE      LENGTH      PLOFFS      STATE      TUTILO      PUTILO
dg rootdg    rootdg     -            -            -            -            -
dm mydg01    c0t1d0s2  -            17678493    -            -            -            -
dm rootdg02  c1t97d0s2 -            17678493    -            -            -            -
dm rootdg03  c1t112d0s2 -            17678493    -            -            -            -
dm rootdg04  c1t114d0s2 -            17678493    -            -            -            -
dm mydg05    c1t96d0s2 -            17678493    -            -            -            -
dm rootdg06  c1t98d0s2 -            17678493    -            -            -            -
dm rootdg07  c1t99d0s2 -            17678493    -            -            -            -
dm rootdg08  c1t100d0s2 -            17678493    -            -            -            -
v  vol1      fsgen      ENABLED      2048        -            ACTIVE      -            -
pl  vol1-01   vol1       ENABLED      3591        -            ACTIVE      -            -
sd mydg01-01 vol1-01    ENABLED      3591        0            -            -            -
pl  vol1-02   vol1       ENABLED      3591        -            ACTIVE      -            -
sd mydg05-01 vol1-02    ENABLED      3591        0            -            -            -
```

See [“Joining shared disk groups”](#) on page 384.

Monitoring and controlling tasks

The VxVM task monitor tracks the progress of system recovery by monitoring task creation, maintenance, and completion. The task monitor lets you monitor task progress and modify characteristics of tasks, such as pausing and recovery rate (for example, to reduce the impact on system performance).

Note: VxVM supports this feature only for private disk groups, not for shared disk groups in a CVM environment.

Specifying task tags

Every task is given a unique task identifier. This is a numeric identifier for the task that can be specified to the `vxtask` utility to specifically identify a single task. Several VxVM utilities also provide a `-t` option to specify an alphanumeric tag of up to 16

characters in length. This allows you to group several tasks by associating them with the same tag.

The following utilities accept the `-t` option:

- `vxassist`
- `vxevac`
- `vxmirror`
- `vxplex`
- `vxrecover`
- `vxrelayout`
- `vxresize`
- `vxsd`
- `vxvol`

For example, to execute a `vxrecover` command and track the resulting tasks as a group with the task tag `myrecovery`, use the following command:

```
# vxrecover -g mydg -t myrecovery -b mydg05
```

To track the resulting tasks, use the following command:

```
# vxtask monitor myrecovery
```

Any tasks started by the utilities invoked by `vxrecover` also inherit its task ID and task tag, establishing a parent-child task relationship.

For more information about the utilities that support task tagging, see their respective manual pages.

Managing tasks with `vxtask`

You can use the `vxtask` command to administer operations on VxVM tasks. Operations include listing tasks, modifying the task state (pausing, resuming, aborting) and modifying the task's progress rate.

VxVM tasks represent long-term operations in progress on the system. Every task gives information on the time the operation started, the size and progress of the operation, and the state and rate of progress of the operation. You can change the state of a task, giving coarse-grained control over the progress of the operation. For those operations that support it, you can change the rate of progress of the task, giving more fine-grained control over the task.

New tasks take time to be set up, and so may not be immediately available for use after a command is invoked. Any script that operates on tasks may need to poll for the existence of a new task.

See the `vxtask(1M)` manual page.

vxtask operations

The `vxtask` command supports the following operations:

<code>abort</code>	Stops the specified task. In most cases, the operations “back out” as if an I/O error occurred, reversing what has been done so far to the largest extent possible.
<code>list</code>	Displays a one-line summary for each task running on the system. The <code>-l</code> option prints tasks in long format. The <code>-h</code> option prints tasks hierarchically, with child tasks following the parent tasks. By default, all tasks running on the system are printed. If you include a <code>taskid</code> argument, the output is limited to those tasks whose <code>taskid</code> or task tag match <code>taskid</code>. The remaining arguments filter tasks and limit which ones are listed. In this release, the <code>vxtask list</code> command supports SmartMove and thin reclamation operation. ■ If you use SmartMove to resync or sync the volume, plex, or subdisk, the <code>vxtask list</code> displays whether the operations is using SmartMove or not. ■ In a LUN level reclamation, the <code>vxtask list</code> command provides information on the amount of the reclaim performed on each LUN. ■ The <code>init=zero</code> on the thin volume may trigger the reclaim on the thin volume and the progress is seen in the <code>vxtask list</code> command.
<code>monitor</code>	Prints information continuously about a task or group of tasks as task information changes. This lets you track task progress. Specifying <code>-l</code> prints a long listing. By default, one-line listings are printed. In addition to printing task information when a task state changes, output is also generated when the task completes. When this occurs, the state of the task is printed as <code>EXITED</code> .
<code>pause</code>	Pauses a running task, causing it to suspend operation.
<code>resume</code>	Causes a paused task to continue operation.

`set` Changes a task's modifiable parameters. Currently, there is only one modifiable parameter, `slow[=iodelay]`, which can be used to reduce the impact that copy operations have on system performance. If you specify `slow`, this introduces a delay between such operations with a default value for `iodelay` of 250 milliseconds. The larger `iodelay` value you specify, the slower the task progresses and the fewer system resources that it consumes in a given time. (The `vxplex`, `vxvol` and `vxrecover` commands also accept the `slow` attribute.)

Using the `vxtask` command

To list all tasks running on the system, use the following command:

```
# vxtask list
```

To print tasks hierarchically, with child tasks following the parent tasks, specify the `-h` option, as follows:

```
# vxtask -h list
```

To trace all paused tasks in the disk group `mydg`, as well as any tasks with the tag `sysstart`, use the following command:

```
# vxtask -g mydg -p -i sysstart list
```

To list all paused tasks, use the `vxtask -p list` command. To continue execution (the task may be specified by its ID or by its tag), use `vxtask resume`:

```
# vxtask -p list
```

```
# vxtask resume 167
```

To monitor all tasks with the tag `myoperation`, use the following command:

```
# vxtask monitor myoperation
```

To cause all tasks tagged with `recoval1` to exit, use the following command:

```
# vxtask abort recoval1
```

This command causes VxVM to try to reverse the progress of the operation so far. For example, aborting an Online Relayout results in VxVM returning the volume to its original layout.

See [“Controlling the progress of a relayout”](#) on page 837.

Using vxnotify to monitor configuration changes

You can use the `vxnotify` utility to display events relating to disk and configuration changes that are managed by the `vxconfigd` configuration daemon. If `vxnotify` is running on a system where the VxVM clustering feature is active, it displays events that are related to changes in the cluster state of the system on which it is running. The `vxnotify` utility displays the requested event types until you kill it, until it has received a specified number of events, or until a specified period of time has elapsed.

Examples of configuration events that can be detected include disabling and enabling of controllers, paths and DMP nodes, RAID-5 volumes entering degraded mode, detachment of disks, plexes and volumes, and nodes joining and leaving a cluster.

For example, the following `vxnotify` command displays information about all disk, plex, and volume detachments as they occur:

```
# vxnotify -f
```

The following command provides information about cluster configuration changes, including the import and deport of shared disk groups:

```
# vxnotify -s -i
```

See the `vxnotify(1M)` manual page.

Performing online relayout

You can use the `vxassist relayout` command to reconfigure the layout of a volume without taking it offline. The general form of this command is as follows:

```
# vxassist [-b] [-g diskgroup] relayout volume [layout=layout] \
    [relayout_options]
```

If you specify the `-b` option, relayout of the volume is a background task.

The following destination layout configurations are supported.

<code>concat-mirror</code>	Concatenated-mirror
<code>concat</code>	Concatenated
<code>nomirror</code>	Concatenated
<code>nostripe</code>	Concatenated
<code>raid5</code>	RAID-5 (not supported for shared disk groups)

span	Concatenated
stripe	Striped

See [“Permitted relayout transformations”](#) on page 832.

For example, the following command changes the concatenated volume `vol102`, in disk group `mydg`, to a striped volume. By default, the striped volume has 2 columns and a 64 KB striped unit size.:

```
# vxassist -g mydg relayout vol102 layout=stripe
```

Sometimes, you may need to perform a relayout on a plex rather than on a volume.

See [“Specifying a plex for relayout”](#) on page 836.

Permitted relayout transformations

[Table 37-1](#) shows the supported relayout transformations for concatenated volumes.

Table 37-1 Supported relayout transformations for concatenated volumes

Relayout to	From concat
concat	No.
concat-mirror	No. Add a mirror, and then use <code>vxassist convert</code> instead.
mirror-concat	No. Add a mirror instead.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes. The stripe width and number of columns may be defined.
stripe	Yes. The stripe width and number of columns may be defined.
stripe-mirror	Yes. The stripe width and number of columns may be defined.

[Table 37-2](#) shows the supported relayout transformations for concatenated-mirror volumes.

Table 37-2 Supported relayout transformations for concatenated-mirror volumes

Relayout to	From concat-mirror
concat	No. Use <code>vxassist convert</code> , and then remove the unwanted mirrors from the resulting mirrored-concatenated volume instead.

Table 37-2 Supported relayout transformations for concatenated-mirror volumes (*continued*)

Relayout to	From concat-mirror
concat-mirror	No.
mirror-concat	No. Use <code>vxassist convert</code> instead.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes.
stripe	Yes. This relayout removes a mirror and adds striping. The stripe width and number of columns may be defined.
stripe-mirror	Yes. The stripe width and number of columns may be defined.

[Table 37-3](#) shows the supported relayout transformations for RAID-5 volumes.

Table 37-3 Supported relayout transformations for mirrored-stripe volumes

Relayout to	From mirror-stripe
concat	Yes.
concat-mirror	Yes.
mirror-concat	No. Use <code>vxassist convert</code> after relayout to the concatenated-mirror volume instead.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes. The stripe width and number of columns may be changed.
stripe	Yes. The stripe width or number of columns must be changed.
stripe-mirror	Yes. The stripe width or number of columns must be changed. Otherwise, use <code>vxassist convert</code> .

[Table 37-4](#) shows the supported relayout transformations for mirror-concatenated volumes.

Table 37-4 Supported relayout transformations for mirrored-concatenated volumes

Relayout to	From mirror-concat
concat	No. Remove the unwanted mirrors instead.
concat-mirror	No. Use <code>vxassist convert</code> instead.
mirror-concat	No.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes. The stripe width and number of columns may be defined. Choose a plex in the existing mirrored volume on which to perform the relayout. The other plexes are removed at the end of the relayout operation.
stripe	Yes.
stripe-mirror	Yes.

[Table 37-5](#) shows the supported relayout transformations for mirrored-stripe volumes.

Table 37-5 Supported relayout transformations for mirrored-stripe volumes

Relayout to	From mirror-stripe
concat	Yes.
concat-mirror	Yes.
mirror-concat	No. Use <code>vxassist convert</code> after relayout to the concatenated-mirror volume instead.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes. The stripe width and number of columns may be changed.
stripe	Yes. The stripe width or number of columns must be changed.
stripe-mirror	Yes. The stripe width or number of columns must be changed. Otherwise, use <code>vxassist convert</code> .

[Table 37-6](#) shows the supported relayout transformations for unmirrored stripe and layered striped-mirror volumes.

Table 37-6 Supported relayout transformations for unmirrored stripe and layered striped-mirror volumes

Relayout to	From stripe or stripe-mirror
concat	Yes.
concat-mirror	Yes.
mirror-concat	No. Use <code>vxassist convert</code> after relayout to the concatenated-mirror volume instead.
mirror-stripe	No. Use <code>vxassist convert</code> after relayout to the striped-mirror volume instead.
raid5	Yes. The stripe width and number of columns may be changed.
stripe	Yes. The stripe width or number of columns must be changed.
stripe-mirror	Yes. The stripe width or number of columns must be changed.

Specifying a non-default layout

You can specify one or more of the following relayout options to change the default layout configuration:

<code>ncol=number</code>	Specifies the number of columns.
<code>ncol+=number</code>	Specifies the number of columns to add.
<code>ncol=-number</code>	Specifies the number of columns to remove.
<code>stripeunit=size</code>	Specifies the stripe width.

The following examples use `vxassist` to change the stripe width and number of columns for a striped volume in the disk group `dbasedg`:

```
# vxassist -g dbasedg relayout vol03 stripeunit=64k ncol=6
# vxassist -g dbasedg relayout vol03 ncol+=2
# vxassist -g dbasedg relayout vol03 stripeunit=128k
```

The following example changes a concatenated volume to a RAID-5 volume with four columns:

```
# vxassist -g dbasedg relayout vol04 layout=raid5 ncol=4
```

Specifying a plex for relayout

If you have enough disks and space in the disk group, you can change any layout to RAID-5. To convert a mirrored volume to RAID-5, you must specify which plex is to be converted. When the conversion finishes, all other plexes are removed, releasing their space for other purposes. If you convert a mirrored volume to a layout other than RAID-5, the unconverted plexes are not removed. Specify the plex to be converted by naming it in place of a volume as follows:

```
# vxassist [-g diskgroup] relayout plex [layout=layout] \  
[relayout_options]
```

Tagging a relayout operation

To control the progress of a relayout operation, for example to pause or reverse it, use the `-t` option to `vxassist` to specify a task tag for the operation. For example, the following relayout is performed as a background task and has the tag `myconv`:

```
# vxassist -b -g dbasedg -t myconv relayout vol04 layout=raid5 \  
ncol=4
```

See [“Viewing the status of a relayout”](#) on page 836.

See [“Controlling the progress of a relayout”](#) on page 837.

Viewing the status of a relayout

Online relayout operations take time to perform. You can use the `vxrelayout` command to obtain information about the status of a relayout operation. For example, the following command:

```
# vxrelayout -g mydg status vol04
```

might display output similar to the following:

```
STRIPED, columns=5, stwidth=128--> STRIPED, columns=6,  
stwidth=128  
Relayout running, 68.58% completed.
```

In this example, the reconfiguration is in progress for a striped volume from 5 to 6 columns, and is over two-thirds complete.

See the `vxrelayout(1M)` manual page.

If you specify a task tag to `vxassist` when you start the relayout, you can use this tag with the `vxtask` command to monitor the progress of the relayout. For example, to monitor the task that is tagged as `myconv`, enter the following:


```
# vxtask monitor myconv
```

Controlling the progress of a relay layout

You can use the `vxtask` command to stop (pause) the relay layout temporarily, or to cancel it (abort). If you specify a task tag to `vxassist` when you start the relay layout, you can use this tag to specify the task to `vxtask`. For example, to pause the relay layout operation that is tagged as `myconv`, enter:

```
# vxtask pause myconv
```

To resume the operation, use the `vxtask` command as follows:

```
# vxtask resume myconv
```

For relay layout operations that have not been stopped using the `vxtask pause` command (for example, the `vxtask abort` command was used to stop the task, the transformation process died, or there was an I/O failure), resume the relay layout by specifying the `start` keyword to `vxrelay layout`, as follows:

```
# vxrelay layout -g mydg -o bg start vol104
```

If you use the `vxrelay layout start` command to restart a relay layout that you previously suspended using the `vxtask pause` command, a new untagged task is created to complete the operation. You cannot then use the original task tag to control the relay layout.

The `-o bg` option restarts the relay layout in the background. You can also specify the `slow` and `iosize` option modifiers to control the speed of the relay layout and the size of each region that is copied. For example, the following command inserts a delay of 1000 milliseconds (1 second) between copying each 10 MB region:

```
# vxrelay layout -g mydg -o bg,slow=1000,iosize=10m start vol104
```

The default delay and region size values are 250 milliseconds and 1 MB respectively.

To reverse the direction of relay layout operation that is stopped, specify the `reverse` keyword to `vxrelay layout` as follows:

```
# vxrelay layout -g mydg -o bg reverse vol104
```

This undoes changes made to the volume so far, and returns it to its original layout.

If you cancel a relay layout using `vxtask abort`, the direction of the conversion is also reversed, and the volume is returned to its original configuration.

See [“Managing tasks with vxtask”](#) on page 828.

See the `vxrelayout(1M)` manual page.

See the `vxtask(1M)` manual page.

Adding a mirror to a volume

You can add a mirror to a volume with the `vxassist` command, as follows:

```
# vxassist [-b] [-g diskgroup] mirror volume
```

Specifying the `-b` option makes synchronizing the new mirror a background task.

For example, to create a mirror of the volume `voltest` in the disk group, `mydg`, use the following command:

```
# vxassist -b -g mydg mirror voltest
```

You can also mirror a volume by creating a plex and then attaching it to a volume using the following commands:

```
# vxmake [-g diskgroup] plex plex sd=subdisk ...  
# vxplex [-g diskgroup] att volume plex
```

Mirroring all volumes

To mirror all volumes in a disk group to available disk space, use the following command:

```
# /etc/vx/bin/vxmirror -g diskgroup -a
```

To configure VxVM to create mirrored volumes by default, use the following command:

```
# vxmirror -d yes
```

If you make this change, you can still make unmirrored volumes by specifying `nmirror=1` as an attribute to the `vxassist` command. For example, to create an unmirrored 20-gigabyte volume named `nomirror` in the disk group `mydg`, use the following command:

```
# vxassist -g mydg make nomirror 20g nmirror=1
```

Mirroring volumes on a VM disk

Mirroring volumes creates one or more copies of your volumes on another disk. By creating mirror copies of your volumes, you protect your volumes against loss of data if a disk fails.

You can use this task on your root disk to make a second copy of the boot information available on an alternate disk. This lets you boot your system even if your root disk fails.

Note: This task only mirrors concatenated volumes. Volumes that are already mirrored or that contain subdisks that reside on multiple disks are ignored

To mirror volumes on a disk

- 1 Make sure that the target disk has an equal or greater amount of space as the source disk.
- 2 From the `vxdiskadm` main menu, select `Mirror volumes on a disk`.
- 3 At the prompt, enter the disk name of the disk that you wish to mirror:

```
Enter disk name [<disk>,list,q,?] mydg02
```

- 4 At the prompt, enter the target disk name (this disk must be the same size or larger than the originating disk):

```
Enter destination disk [<disk>,list,q,?] (default: any) mydg01
```

- 5 At the prompt, press **Return** to make the mirror:

```
Continue with operation? [y,n,q,?] (default: y)
```

The `vxdiskadm` program displays the status of the mirroring operation, as follows:

```
VxVM vxmirror INFO V-5-2-22 Mirror volume voltest-bk00
.
.
.
VxVM INFO V-5-2-674 Mirroring of disk mydg01 is complete.
```

- 6 At the prompt, indicate whether you want to mirror volumes on another disk (y) or return to the `vxdiskadm` main menu (n):

```
Mirror volumes on another disk? [y,n,q,?] (default: n)
```

Additional mirroring considerations

The larger private region size that was introduced in VxVM 3.2 (1MB) and VxVM 5.0 (32MB) may create one of the following mirroring scenarios under which `vxdiskadm` fails:

- Mirroring a full root disk to a target disk that is the same size as the source disk. A full disk has no free cylinders.
- Mirroring a disk created using an earlier version of Veritas Volume Manager to a target disk that is the same size as the source disk. You only need to use this step if mirroring using `vxdiskadm` fails.
- Mirroring a full Veritas Volume Manager disk (not a root disk) that was encapsulated in VxVM 3.5 to a target disk that is the same size as the source disk. You only need to use this step if mirroring using `vxdiskadm` fails.

See the `vxdiskadm(1M)` manual page.

To create a mirror under any of these scenarios

- 1 Determine the size of the source disk's private region, using one of the following methods:

- If the source disk is a root disk, obtain its private region length by running the following command:

```
# vxprint -l rootdisk
```

The disk media name of the root disk is typically `rootdisk`.

In the output, find the `privlen` value. In this example, the value is 3071:

```
devinfo: publen=39846240 privlen=3071
```

- If the source disk is not a root disk, obtain its private region length by running the following command:

```
# vxdisk list diskname
```

where `diskname` is the disk media name of the source disk.

In the displayed output, note the `len` value for the `private` field. In this example, the value of this field is 3071:

```
private: slice=4 offset=1 len=3071
```

2 Use the `vxdisksetup` program to initialize the target disk, Enter the following:

```
# /usr/lib/vxvm/bin/vxdisksetup -i c#t#d# privoffset=0 \
  privlen=XXXX publen=YYYY
```

where `XXXX` is the size of the source disk’s private region, and `YYYY` is the size of its public region.

If your system is configured to use enclosure-based naming instead of OS-based naming, replace the `c#t#d#` name with the enclosure-based name for the disk.

3 Add the newly initialized target disk to the source disk group. Enter the following:

```
# vxdg -g diskgroup adddisk medianame=c#t#d#
```

4 Use the `vxdiskadm` command and select `Mirror` volumes on a disk to create the mirror. Specify the disk media names of the source disk (*rootdisk*) and the target disk (*medianame*).

Configuring SmartMove

By default, the SmartMove utility is enabled for all volumes. Configuring the SmartMove feature is only required if you want to change the default behavior, or if you have modified the behavior previously.

SmartMove has three values where SmartMove can be applied or not. The three values are:

Value	Meaning
<code>none</code>	Do not use SmartMove at all.
<code>thinonly</code>	Use SmartMove for thin aware LUNs only.
<code>all</code>	Use SmartMove for all types of LUNs.
	This is the default value.

To configure the SmartMove value

- 1 To display the current and default SmartMove values, type the following command:

```
# vxdefault list
KEYWORD                                CURRENT-VALUE  DEFAULT-VALUE
usefssmartmove                         all            all
...
```

- 2 To set the SmartMove value, type the following command:

```
# vxdefault set usefssmartmove value
```

where *value* is either *none*, *thinonly*, or *all*.

Removing a mirror

When you no longer need a mirror, you can remove it to free disk space.

Note: VxVM will not allow you to remove the last valid plex associated with a volume.

To remove a mirror from a volume, use the following command:

```
# vxassist [-g diskgroup] remove mirror volume
```

You can also use storage attributes to specify the storage to be removed. For example, to remove a mirror on disk *mydg01* from volume *vol01*, enter the following.

Note: The **!** character is a special character in some shells. The following example shows how to escape it in a bash shell.

```
# vxassist -g mydg remove mirror vol01 \!mydg01
```

See [“Creating a volume on specific disks”](#) on page 204.

Alternatively, use the following command to dissociate and remove a mirror from a volume:

```
# vxplex [-g diskgroup] -o rm dis mirror
```

For example, to dissociate and remove a mirror named *vol01-02* from the disk group *mydg*, use the following command:

```
# vxplex -g mydg -o rm dis vol01-02
```

This command removes the mirror `vol101-02` and all associated subdisks. This is equivalent to entering the following commands separately:

```
# vxplex -g mydg dis vol101-02
# vxedit -g mydg -r rm vol101-02
```

Setting tags on volumes

Volume tags implement the SmartTier feature. You can also apply tags to vsets using the same `vxvm` command syntax as shown below.

The following forms of the `vxassist` command let you do the following:

- Set a named tag and optional tag value on a volume.
- Replace a tag.
- Remove a tag from a volume.

```
# vxassist [-g diskgroup] settag volume|vset tagname[=tagvalue]
# vxassist [-g diskgroup] replacetag volume|vset oldtag newtag
# vxassist [-g diskgroup] removetag volume|vset tagname
```

To list the tags that are associated with a volume, use the following command:

```
# vxassist [-g diskgroup] listtag [volume|vset]
```

If you do not specify a volume name, all the volumes and vsets in the disk group are displayed. The acronym `vt` in the `TY` field indicates a vset.

The following is a sample `listtag` command:

```
# vxassist -g dgl listtag vol
```

To list the volumes that have a specified tag name, use the following command:

```
# vxassist [-g diskgroup] list tag=tagname volume
```

Tag names and tag values are case-sensitive character strings of up to 256 characters. Tag names can consist of the following ASCII characters:

- Letters (A through Z and a through z)
- Numbers (0 through 9)
- Dashes (-)
- Underscores (_)
- Periods (.)

A tag name must start with either a letter or an underscore.

Tag values can consist of any ASCII character that has a decimal value from 32 through 127. If a tag value includes spaces, quote the specification to protect it from the shell, as follows:

```
# vxassist -g mydg settag myvol "dbvol=table space 1"
```

The `list` operation understands dotted tag hierarchies. For example, the listing for `tag=a.b` includes all volumes that have tag names starting with `a.b`.

The tag names `site`, `udid`, and `vdid` are reserved. Do not use them. To avoid possible clashes with future product features, do not start tag names with any of the following strings: `asl`, `be`, `nbu`, `sf`, `symc`, or `vx`.

Managing disk groups

This section describes managing disk groups.

Disk group versions

All disk groups have a version number associated with them. Each major Veritas Volume Manager (VxVM) release introduces a disk group version. To support the new features in the release, the disk group must be the latest disk group version. By default, VxVM creates disk groups with the latest disk group version. For example, Veritas Volume Manager 6.0 creates disk groups with version 170.

Each VxVM release supports a specific set of disk group versions. VxVM can import and perform operations on a disk group of any supported version. However, the operations are limited by what features and operations the disk group version supports. If you import a disk group from a previous version, the latest features may not be available. If you attempt to use a feature from a newer version of VxVM, you receive an error message similar to this:

```
VxVM vxedit ERROR V-5-1-2829 Disk group version doesn't support  
feature
```

You must explicitly upgrade the disk group to the appropriate disk group version to use the feature.

See [“Upgrading the disk group version”](#) on page 848.

[Table 37-7](#) summarizes the Veritas Volume Manager releases that introduce and support specific disk group versions. It also summarizes the features that are supported by each disk group version.

Table 37-7 Disk group version assignments

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
6.0	170	■ VVR compression ■ VVR Secondary logging ■ CVM availability enhancements ■ DCO version 30 ■ Recovery for synchronization tasks.	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130, 140, 150, 160
5.1SP1	160	■ Automated bunker replay as part of GCO failover ■ Ability to elect primary during GCO takeover ■ CVM support for more than 32 nodes and up to 64 nodes ■ CDS layout support for large luns (> 1 TB) ■ vxrootadm enhancements	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130, 140, 150, 160
5.1	150	SSD device support, migration of ISP dg	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130, 140, 150
5.0	140	Data migration, Remote Mirror, coordinator disk groups (used by VCS), linked volumes, snapshot LUN import.	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130, 140
5.0	130	■ VVR Enhancements	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130

Table 37-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
4.1	120	- Automatic Cluster-wide Failback for A/P arrays - Persistent DMP Policies - Shared Disk Group Failure Policy	20, 30, 40, 50, 60, 70, 80, 90, 110, 120
4.0	110	- Cross-platform Data Sharing (CDS) - Device Discovery Layer (DDL) 2.0 - Disk Group Configuration Backup and Restore - Elimination of <code>rootdg</code> as a Special Disk Group - Full-Sized and Space-Optimized Instant Snapshots - Intelligent Storage Provisioning (ISP) - Serial Split Brain Detection - Volume Sets (Multiple Device Support for VxFS)	20, 30, 40, 50, 60, 70, 80, 90, 110

Table 37-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
3.2, 3.5	90	- Cluster Support for Oracle Resilvering - Disk Group Move, Split and Join - Device Discovery Layer (DDL) 1.0 - Layered Volume Support in Clusters - Ordered Allocation - OS Independent Naming Support - Persistent FastResync	20, 30, 40, 50, 60, 70, 80, 90
3.1.1	80	VVR Enhancements	20, 30, 40, 50, 60, 70, 80
3.1	70	- Non-Persistent FastResync - Sequential DRL - Unrelocate - VVR Enhancements	20, 30, 40, 50, 60, 70
3.0	60	- Online Relayout - Safe RAID-5 Subdisk Moves	20, 30, 40, 60
2.5	50	SRVM (now known as Veritas Volume Replicator or VVR)	20, 30, 40, 50
2.3	40	Hot-Relocation	20, 30, 40
2.2	30	VxSmartSync Recovery Accelerator	20, 30

Table 37-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
2.0	20	■ Dirty Region Logging (DRL) ■ Disk Group Configuration Copy Limiting ■ Mirrored Volumes Logging ■ New-Style Stripes ■ RAID-5 Volumes ■ Recovery Checkpointing	20
1.3	15		15
1.2	10		10

If you need to import a disk group on a system running an older version of Veritas Volume Manager, you can create a disk group with an earlier disk group version.

See [“Creating a disk group with an earlier disk group version”](#) on page 849.

Upgrading the disk group version

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version.

When you upgrade, VxVM does not automatically upgrade the versions of existing disk groups. If the disk group is a supported version, the disk group can be used “as is”, as long as you do not attempt to use the features of the current version. Until the disk group is upgraded, it may still be deported back to the release from which it was imported.

To use the features in the upgraded release, you must explicitly upgrade the existing disk groups. There is no “downgrade” facility. After you upgrade a disk group, the disk group is incompatible with earlier releases of VxVM that do not support the new version. For disk groups that are shared among multiple servers for failover or for off-host processing, verify that the VxVM release on all potential hosts that may use the disk group supports the disk group version to which you are upgrading.

After upgrading to Storage Foundation Cluster File System High Availability , you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

To list the version of a disk group, use this command:

```
# vxdg list dgname
```

You can also determine the disk group version by using the `vxprint` command with the `-l` format option.

To upgrade a disk group to the highest version supported by the release of VxVM that is currently running, use this command:

```
# vxdg upgrade dgname
```

Creating a disk group with an earlier disk group version

You may sometimes need to create a disk group that can be imported on a system running an older version of Veritas Volume Manager. You must specify the disk group version when you create the disk group, since you cannot downgrade a disk group version.

For example, the default disk group version for a disk group created on a system running Veritas Volume Manager 6.0 is 170. Such a disk group cannot be imported on a system running Veritas Volume Manager 4.1, as that release only supports up to version 120. Therefore, to create a disk group on a system running Veritas Volume Manager 6.0 that can be imported by a system running Veritas Volume Manager 4.1, the disk group must be created with a version of 120 or less.

To create a disk group with a previous version, specify the `-T version` option to the `vxdg init` command.

For example, to create a disk group with version 120 that can be imported by a system running VxVM 4.1, use the following command:

```
# vxdg -T 120 init newdg newdg01=c0t3d0s2
```

This creates a disk group, `newdg`, which can be imported by Veritas Volume Manager 4.1. Note that while this disk group can be imported on the VxVM 4.1 system, attempts to use features from Veritas Volume Manager 5.0 and later releases will fail.

Managing volumes

This section describes managing volumes.

Managing plexes and subdisks

This section describes managing plexes and subdisks.

A subdisk is a set of contiguous disk blocks. VxVM allocates disk space using subdisks.

A plex is a logical groupings of subdisks that creates an area of disk space independent of physical disk size or other restrictions. Replication (mirroring) of disk data is set up by creating multiple data plexes for a single volume. Each data plex in a mirrored volume contains an identical copy of the volume data.

A plex becomes a participating plex for a volume when it is attached to a volume. Attaching a plex associates it with the volume and enables the plex for use.

Reattaching plexes

When a mirror plex encounters irrecoverable errors, Veritas Volume Manager (VxVM) detaches the plex from the mirrored volume. An administrator may also detach a plex manually using a utility such as `vxplex` or `vxassist`. In order to use a plex that was previously attached to a volume, the plex must be reattached to the volume. The reattach operation also ensures that the plex mirror is resynchronized to the other plexes in the volume.

See [“Plex synchronization”](#) on page 853.

The following methods are available for reattaching plexes:

- By default, VxVM automatically reattaches the affected mirror plexes when the underlying failed disk or LUN becomes visible. When VxVM detects that the device is online, VxVM automatically recovers the volume components on the involved LUN. VxVM resynchronizes the plex and the mirror becomes available. See [“Automatic plex reattachment”](#) on page 850.
- If the automatic reattachment feature is disabled, you need to reattach the plexes manually. You may also need to manually reattach the plexes for devices that are not automatically reattached. For example, VxVM does not automatically reattach plexes on site-consistent volumes. See [“Reattaching a plex manually”](#) on page 852.

Automatic plex reattachment

When a mirror plex encounters irrecoverable errors, Veritas Volume Manager (VxVM) detaches the plex from the mirrored volume. By default, VxVM automatically reattaches the affected mirror plexes when the underlying failed disk or LUN becomes visible. When VxVM detects that the device is online, the VxVM volume

components on the involved LUN are automatically recovered, and the mirrors become usable.

VxVM uses the DMP failed LUN probing to detect when the device has come online. The timing for a reattach depends on the `dmp_restore_interval`, which is a tunable parameter. The number of LUNs that have reconnected may also affect the time required before the plex is reattached.

VxVM does not automatically reattach plexes on site-consistent volumes.

When VxVM is installed or the system reboots, VxVM starts the `vxattachd` daemon. The `vxattachd` daemon handles automatic reattachment for both plexes and sites. The `vxattachd` daemon also initiates the resynchronization process for a plex. After a plex is successfully reattached, `vxattachd` notifies root.

To disable automatic plex attachment, remove `vxattachd` from the start up scripts. Disabling `vxattachd` disables the automatic reattachment feature for both plexes and sites.

In a Cluster Volume Manager (CVM) the following considerations apply:

- If the global detach policy is set, a storage failure from any node causes all plexes on that storage to be detached globally. When the storage is connected back to any node, the `vxattachd` daemon triggers reattaching the plexes on the master node only.
- The automatic reattachment functionality is local to a node. When enabled on a node, all of the disk groups imported on the node are monitored. If the automatic reattachment functionality is disabled on a master node, the feature is disabled on all shared disk groups and private disk groups imported on the master node.
- The `vxattachd` daemon listens for "dmpnode online" events using `vxnotify` to trigger its operation. Therefore, an automatic reattachment is not triggered if the `dmpnode` online event is not generated when `vxattachd` is running. The following are typical examples:
 - Storage is reconnected before `vxattachd` is started; for example, during reboot.
 - In CVM, with active/passive arrays, if all nodes cannot agree on a common path to an array controller, a plex can get detached due to I/O failure. In these cases, the `dmpnode` will not get disabled. Therefore, after the connections are restored, a `dmpnode` online event is not generated and automatic plex reattachment is not triggered.

These CVM considerations also apply to automatic site reattachment.

See [“Automatic site reattachment”](#) on page 482.

Reattaching a plex manually

This section describes how to reattach plexes manually if automatic reattachment feature is disabled. This procedure may also be required for devices that are not automatically reattached. For example, VxVM does not automatically reattach plexes on site-consistent volumes.

When a disk has been repaired or replaced and is again ready for use, the plexes must be put back online (plex state set to `ACTIVE`). To set the plexes to `ACTIVE`, use one of the following procedures depending on the state of the volume.

- If the volume is currently `ENABLED`, use the following command to reattach the plex:

```
# vxplex [-g diskgroup] att volume plex ...
```

For example, for a plex named `vol101-02` on a volume named `vol101` in the disk group, `mydg`, use the following command:

```
# vxplex -g mydg att vol101 vol101-02
```

As when returning an `OFFLINE` plex to `ACTIVE`, this command starts to recover the contents of the plex and, after the recovery is complete, sets the plex utility state to `ACTIVE`.

- If the volume is not in use (not `ENABLED`), use the following command to re-enable the plex for use:

```
# vxmend [-g diskgroup] on plex
```

For example, to re-enable a plex named `vol101-02` in the disk group, `mydg`, enter:

```
# vxmend -g mydg on vol101-02
```

In this case, the state of `vol101-02` is set to `STALE`. When the volume is next started, the data on the plex is revived from another plex, and incorporated into the volume with its state set to `ACTIVE`.

If the `vxinfo` command shows that the volume is unstartable, set one of the plexes to `CLEAN` using the following command:

```
# vxmend [-g diskgroup] fix clean plex
```

Start the volume using the following command:

```
# vxvol [-g diskgroup] start volume
```

See the *Veritas Storage Foundation High Availability Solutions Troubleshooting Guide*.

Plex synchronization

Each plex or mirror of a volume is a complete copy of the data. When a plex is attached to a volume, the data in the plex must be synchronized with the data in the other plexes in the volume. The plex that is attached may be a new mirror or a formerly attached plex. A new mirror must be fully synchronized. A formerly attached plex only requires the changes that were applied since the plex was detached.

The following operations trigger a plex synchronization:

- Moving or copying a subdisk with the `vxsd` command. The operation creates a temporary plex that is synchronized with the original subdisk.
- Adding a mirror with the `vxassist mirror` command.
- Creating a volume with a mirror with the `vxassist make` command.
- Manually reattaching a plex with the `vxplex att` command.
- Recovering a volume with the `vxrecover` command.
- Adding a mirror to a snapshot with the `vxsnap addmir` command.
- Reattaching or restoring a snapshot with the `vxsnap` command.

Plex synchronization can be a long-running operation, depending on the size of the volume and the amount of data that needs to be synchronized. Veritas Volume Manager provides several features to improve the efficiency of synchronizing the plexes.

- **FastResync**

If the FastResync feature is enabled, VxVM maintains a FastResync map on the volume. VxVM uses the FastResync map to apply only the updates that the mirror has missed. This behavior provides an efficient way to resynchronize the plexes.

- **SmartMove**

The SmartMove™ feature reduces the time and I/O required to attach or reattach a plex to a VxVM volume with a mounted VxFS file system. The SmartMove feature uses the VxFS information to detect free extents and avoid copying them.

When the SmartMove feature is on, less I/O is sent through the host, through the storage network and to the disks or LUNs. The SmartMove feature can be used for faster plex creation and faster array migrations.

- **Recovery for synchronization tasks**

In this release, VxVM tracks the plex synchronization for the following commands: `vxplex att`, `vxassist mirror`, `vxsnap addmir`, `vxsnap reattach`, and `vxsnap restore`. If the system crashes or the `vxconfigd` daemon fails, VxVM provides automatic recovery for the synchronization task. When the system is recovered,

VxVM restarts the synchronization from the point where it failed. The synchronization occurs in the background, so the volume is available without delay.

Decommissioning storage

This section describes how you remove disks and volumes from VxVM.

Removing a volume

If a volume is inactive or its contents have been archived, you may no longer need it. In that case, you can remove the volume and free up the disk space for other uses.

To remove a volume

- 1 Remove all references to the volume by application programs, including shells, that are running on the system.
- 2 If the volume is mounted as a file system, unmount it with the following command:

```
# umount /dev/vx/dsk/diskgroup/volume
```

- 3 If the volume is listed in the `/etc/vfstab` file, edit this file and remove its entry. For more information about the format of this file and how you can modify it, see your operating system documentation
- 4 Stop all activity by VxVM on the volume with the following command:

```
# vxvol [-g diskgroup] stop volume
```

- 5 Remove the volume using the `vxassist` command as follows:

```
# vxassist [-g diskgroup] remove volume volume
```

You can also use the `vxedit` command to remove the volume as follows:

```
# vxedit [-g diskgroup] [-r] [-f] rm volume
```

The `-r` option to `vxedit` indicates recursive removal. This command removes all the plexes that are associated with the volume and all subdisks that are associated with the plexes. The `-f` option to `vxedit` forces removal. If the volume is still enabled, you must specify this option.

Removing a disk from VxVM control

After removing a disk from a disk group, you can permanently remove it from Veritas Volume Manager control.

Warning: The `vxdiskunsetup` command removes a disk from Veritas Volume Manager control by erasing the VxVM metadata on the disk. To prevent data loss, any data on the disk should first be evacuated from the disk. The `vxdiskunsetup` command should only be used by a system administrator who is trained and knowledgeable about Veritas Volume Manager.

To remove a disk from VxVM control

- ◆ Type the following command:

```
# /usr/lib/vxvm/bin/vxdiskunsetup c#t#d#
```

See the `vxdiskunsetup(1m)` manual page.

About shredding data

When you decommission a disk that contained sensitive data, you may need to destroy any remaining data on the disk. Simply deleting the data may not adequately protect the confidential and secure data. In addition to deleting the data, you want to prevent the possibility that hackers can recover any information that is stored on the disks. Regulatory standards require that the confidential and secure data is sanitized or erased using a method such as overwriting the data with a digital pattern. Veritas Volume Manager (VxVM) provides the disk shred operation, which overwrites all of the addressable blocks with a digital pattern in one, three, or seven passes.

Caution: All data in the volume will be lost when you shred it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

VxVM provides the ability to shred the data on the disk to minimize the chance that the data is recoverable. When you specify the disk shred operation, VxVM shreds the entire disk, including any existing disk labels. After the shred operation, VxVM writes a new empty label on the disk to prevent the disk from going to the error state. The VxVM shred operation provides the following methods of overwriting a disk:

- One-pass algorithm

VxVM overwrites the disk with a randomly-selected digital pattern. This option takes the least amount of time. The default type is the one-pass algorithm.

- **Three-pass algorithm**
VxVM overwrites the disk a total of three times. In the first pass, VxVM overwrites the data with a pre-selected digital pattern. The second time, VxVM overwrites the data with the binary complement of the pattern. In the last pass, VxVM overwrites the disk with a randomly-selected digital pattern.
- **Seven-pass algorithm**
VxVM overwrites the disk a total of seven times. In each pass, VxVM overwrites the data with a randomly-selected digital pattern or with the binary complement of the previous pattern.

VxVM does not currently support shredding of thin-reclaimable LUNs. If you attempt to start the shred operation on a thin-reclaimable disk, VxVM displays a warning message and skips the disk.

Shredding a VxVM disk

When you decommission a Veritas Volume Manager (VxVM) disk that contains sensitive data, VxVM provides the ability to shred the data on the disk.

Note the following requirements:

- VxVM does not shred a disk that is in use by VxVM on this system or in a shared disk group.
- VxVM does not currently support shredding of thin-reclaimable LUNs. If you attempt to start the shred operation on a thin-reclaimable disk, VxVM displays a warning message and skips the disk.
- VxVM does not shred a disk that is not a VxVM disk.
- VxVM does not shred a disk that is mounted.
- Symantec does not recommend shredding solid state drives (SSDs). To shred SSD devices, use the shred operation with the force (-f) option.

See [“About shredding data”](#) on page 855.

Caution: All data on the disk will be lost when you shred the disk. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

To shred a VxVM disk

1 To shred the disk:

```
# /etc/vx/bin/vxdiskunsetup [-Cf] -o shred[=1|3|7] disk...
```

Where:

The force option (-f) permits you to shred Solid State Drives (SSDs).

1, 3 and 7 are the shred options corresponding to the number of passes. The default number of passes is 1.

disk... represents one or more disk names. If you specify multiple disk names, the `vxdiskunsetup` command processes them sequentially, one at a time.

For example:

```
# /etc/vx/bin/vxdiskunsetup -o shred=3 hds9970v0_14
disk_shred: Shredding disk hds9970v0_14 with type 3
disk_shred: Disk raw size 2097807360 bytes
disk_shred: Writing 32010 (65536 byte size) pages and 0 bytes
to disk
disk_shred: Wipe Pass 0: Pattern 0x3e
disk_shred: Wipe Pass 1: Pattern 0xca
disk_shred: Wipe Pass 2: Pattern 0xe2
disk_shred: Shred passed random verify of 131072 bytes at
offset 160903168
```

The `vxdiskunsetup shred` command sets up a new task.

2 You can monitor the progress of the shred operation with the `vxtask` command.

For example:

```
# vxtask list
TASKID PTID TYPE/STATE PCT PROGRESS
203 - DISKSHRED/R 90.16% 0/12291840/11081728 DISKSHRED
nodg nodg
```

You can pause, abort, or resume the shred task. You cannot throttle the shred task.

See `vxtask(1m)`

3 If the disk shred operation fails, the disk may go into an error state with no label.

See [“Failed disk shred operation results in a disk with no label”](#) on page 858.

Failed disk shred operation results in a disk with no label

The disk shred operation destroys the label for the disk and recreates the label. If the shred operation aborts in the middle or the system crashes, the disk might go in an error state with no label.

To correct the error state of the disk

- 1 Create a new label manually or reinitialize the disk under VxVM using the following command:

```
# /etc/vx/bin/vxdisksetup -i disk
```

- 2 Start the shred operation. If the disk shows as a non-VxVM disk, reinitialize the disk with the vxdisksetup command in step 1, then restart the shred operation.

```
# /etc/vx/bin/vxdiskunsetup [-Cf] -o shred[=1|3|7] disk...
```

Removing and replacing disks

A replacement disk should have the same disk geometry as the disk that failed. That is, the replacement disk should have the same bytes per sector, sectors per track, tracks per cylinder and sectors per cylinder, same number of cylinders, and the same number of accessible cylinders.

You can use the `prtvtoc` command to obtain disk information.

Note: You may need to run commands that are specific to the operating system or disk array before removing a physical disk.

The removal and replacement of a disk in a Sun StorEdge™ A5x00 or similar type of array, requires a different procedure.

See [“Removing and replacing a disk in a Sun StorEdge A5x00 disk array”](#) on page 864.

If failures are starting to occur on a disk, but the disk has not yet failed completely, you can replace the disk. This involves detaching the failed or failing disk from its disk group, followed by replacing the failed or failing disk with a new one. Replacing the disk can be postponed until a later date if necessary.

If removing a disk causes a volume to be disabled, you can restart the volume so that you can restore its data from a backup.

See the *Storage Foundation High Availability Solutions Troubleshooting Guide*.

To replace a disk

- 1 Select Remove a disk for replacement from the `vxdiskadm` main menu.
- 2 At the following prompt, enter the name of the disk to be replaced (or enter `list` for a list of disks):

```
Enter disk name [<disk>,list,q,?] mydg02
```

- 3 When you select a disk to remove for replacement, all volumes that are affected by the operation are displayed, for example:

```
VxVM NOTICE V-5-2-371 The following volumes will lose mirrors  
as a result of this operation:
```

```
home src
```

```
No data on these volumes will be lost.
```

```
The following volumes are in use, and will be disabled as a  
result of this operation:
```

```
mkting
```

```
Any applications using these volumes will fail future  
accesses. These volumes will require restoration from backup.
```

```
Are you sure you want do this? [y,n,q,?] (default: n)
```

To remove the disk, causing the named volumes to be disabled and data to be lost when the disk is replaced, enter `y` or press Return.

To abandon removal of the disk, and back up or move the data associated with the volumes that would otherwise be disabled, enter `n` or `q` and press Return.

For example, to move the volume `mkting` to a disk other than `mydg02`, use the following command.

The `!` character is a special character in some shells. The following example shows how to escape it in a bash shell.

```
# vxassist move mkting \!mydg02
```

After backing up or moving the data in the volumes, start again from step 1.

- 4 At the following prompt, either select the device name of the replacement disk (from the list provided), press Return to choose the default disk, or enter `none` if you are going to replace the physical disk:

```
The following devices are available as replacements:  
c0t1d0
```

```
You can choose one of these disks now, to replace mydg02.  
Select none if you do not wish to select a replacement disk.
```

```
Choose a device, or select none  
[<device>,none,q,?] (default: c0t1d0)
```

Do not choose the old disk drive as a replacement even though it appears in the selection list. If necessary, you can choose to initialize a new disk.

You can enter `none` if you intend to replace the physical disk.

See [“Replacing a failed or removed disk”](#) on page 862.

- 5 If you chose to replace the disk in step 4, press Return at the following prompt to confirm this:

```
VxVM NOTICE V-5-2-285 Requested operation is to remove mydg02  
from group mydg. The removed disk will be replaced with disk device  
c0t1d0. Continue with operation? [y,n,q,?] (default: y)
```

`vxdiskadm` displays the following messages to indicate that the original disk is being removed:

```
VxVM NOTICE V-5-2-265 Removal of disk mydg02 completed  
successfully.  
VxVM NOTICE V-5-2-260 Proceeding to replace mydg02 with device  
c0t1d0.
```


- 6 If the disk was previously an encapsulated root disk, `vxdiskadm` displays the following message. Enter `y` to confirm that you want to reinitialize the disk:

```
The disk clt0d0 was a previously encapsulated root disk. Due
to the disk layout that results from root disk encapsulation,
the preferred action is to reinitialize and reorganize this
disk. However, if you have any non-redundant data on this disk
you should not reorganize this disk, as the data will be lost.
Reorganize the disk? [y,n,q,?] (default: y) y
```

Warning: It is recommended that you do not enter `n` at this prompt. This results in an invalid `vTOC` that makes the disk unbootable.

Entering `y` at the prompt destroys any data that is on the disk. Ensure that you have at least one valid copy of the data on other disks before proceeding.

- 7 You can now choose whether the disk is to be formatted as a CDS disk that is portable between different operating systems, or as a non-portable sliced or simple disk:

```
Enter the desired format [cdsdisk,sliced,simple,q,?]
(default: cdsdisk)
```

Enter the format that is appropriate for your needs. In most cases, this is the default format, `cdsdisk`.

- 8 At the following prompt, `vxdiskadm` asks if you want to use the default private region size of 65536 blocks (32 MB). Press Return to confirm that you want to use the default value, or enter a different value. (The maximum value that you can specify is 524288 blocks.)

```
Enter desired private region length [<privlen>,q,?]
(default: 65536)
```

- 9 If one of more mirror plexes were moved from the disk, you are now prompted whether FastResync should be used to resynchronize the plexes:

```
Use FMR for plex resync? [y,n,q,?] (default: n) y
vxdiskadm displays the following success message:
VxVM NOTICE V-5-2-158 Disk replacement completed successfully.
```

- 10 At the following prompt, indicate whether you want to remove another disk (y) or return to the vxdiskadm main menu (n):

```
Remove another disk? [y,n,q,?] (default: n)
```

It is possible to move hot-relocate subdisks back to a replacement disk.

See [“Configuring hot-relocation to use only spare disks”](#) on page 782.

Replacing a failed or removed disk

The following procedure describes how to replace a failed or removed disk.

A different procedure is required to remove and replace a disk in a Sun StorEdge A5x00 or similar type of array.

See [“Removing and replacing a disk in a Sun StorEdge A5x00 disk array”](#) on page 864.

To specify a disk that has replaced a failed or removed disk

- 1 Select **Replace a failed or removed disk** from the vxdiskadm main menu.
- 2 At the following prompt, enter the name of the disk to be replaced (or enter list for a list of disks):

```
Select a removed or failed disk [<disk>,list,q,?] mydg02
```

- 3 The vxdiskadm program displays the device names of the disk devices available for use as replacement disks. Your system may use a device name that differs from the examples. Enter the device name of the disk or press Return to select the default device:

```
The following devices are available as replacements:
c0t1d0 c1t1d0
```

```
You can choose one of these disks to replace mydg02.
Choose "none" to initialize another disk to replace mydg02.
```

```
Choose a device, or select "none"
[<device>,none,q,?] (default: c0t1d0)
```

4 Depending on whether the replacement disk was previously initialized, perform the appropriate step from the following:

- If the disk has not previously been initialized, press Return at the following prompt to replace the disk:

```
VxVM INFO V-5-2-378 The requested operation is to initialize
disk device c0t1d0 and to then use that device to
replace the removed or failed disk mydg02 in disk group mydg.
Continue with operation? [y,n,q,?] (default: y)
```

- If the disk has already been initialized, press Return at the following prompt to replace the disk:

```
VxVM INFO V-5-2-382 The requested operation is to use the
initialized device c0t1d0 to replace the removed or
failed disk mydg02 in disk group mydg.
Continue with operation? [y,n,q,?] (default: y)
```

- If the disk was previously an encapsulated root disk, `vxdiskadm` displays the following message. Enter `y` to confirm that you want to reinitialize the disk:

```
VxVM INFO V-5-2-876 The disk c0t1d0 was a previously
encapsulated root disk. Due to the disk layout that results
from root disk encapsulation, the preferred action is to
reinitialize and reorganize this disk. However, if you have
any non-redundant data on this disk you should not reorganize
this disk, as the data will be lost.
Reorganize the disk? [y,n,q,?] (default: y) y
```

Warning: It is recommended that you do not enter `n` at this prompt. This can result in an invalid VTOC that makes the disk unbootable.

Entering `y` at the prompt destroys any data that is on the disk. Ensure that you have at least one valid copy of the data on other disks before proceeding.

- 5 You can now choose whether the disk is to be formatted as a CDS disk that is portable between different operating systems, or as a non-portable sliced or simple disk:

```
Enter the desired format [cdsdisk,sliced,simple,q,?]
(default: cdsdisk)
```

Enter the format that is appropriate for your needs. In most cases, this is the default format, `cdsdisk`.

- 6 At the following prompt, `vxdiskadm` asks if you want to use the default private region size of 65536 blocks (32 MB). Press Return to confirm that you want to use the default value, or enter a different value. (The maximum value that you can specify is 524288 blocks.)

```
Enter desired private region length [<privlen>,q,?]
(default: 65536)
```

- 7 The `vxdiskadm` program then proceeds to replace the disk, and returns the following message on success:

```
VxVM NOTICE V-5-2-158 Disk replacement completed successfully.
```

At the following prompt, indicate whether you want to replace another disk (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Replace another disk? [y,n,q,?] (default: n)
```

Removing and replacing a disk in a Sun StorEdge A5x00 disk array

Note: The following procedure is suitable for use with any array that is administered by using the Solaris `luxadm` command.

To replace a disk in a Sun StorEdge A5x00 disk array

- 1 Run the `vxdiskadm` command, and select `Remove a disk for replacement` from the main menu. Enter `none` when prompted to name a replacement disk.
- 2 Use the following command to remove the disk from VxVM:

```
# vxdisk rm daname
```

where *daname* is the disk access name of the device (for example, `c1t5d0s2`).

- 3 Use the Solaris `luxadm` command to obtain the array name and slot number of the disk, and then use these values with `luxadm` to remove the disk:

```
# luxadm disp /dev/rdisk/daname  
# luxadm remove_device array_name,slot_number
```

Follow the `luxadm` prompts, and pull out the disk when instructed.

- 4 Run the following `luxadm` command when you are ready to insert the replacement disk:

```
# luxadm insert_device array_name,slot_number
```

Follow the `luxadm` prompts, and insert the replacement disk when instructed.

- 5 Run the following command to scan for the new disk and update the system:

```
# vxdiskconfig
```

- 6 Run the `vxdiskadm` command, select `Replace a failed or removed disk` from the main menu, and follow the instructions.

See [“Replacing a failed or removed disk”](#) on page 862.

Rootability

This chapter includes the following topics:

- [Encapsulating a disk](#)
- [Rootability](#)
- [Administering an encapsulated boot disk](#)
- [Unencapsulating the root disk](#)

Encapsulating a disk

Warning: Encapsulating a disk requires that the system be rebooted several times. Schedule performance of this procedure for a time when this does not inconvenience users.

This section describes how to encapsulate a disk for use in VxVM. Encapsulation preserves any existing data on the disk when the disk is placed under VxVM control.

To prevent the encapsulation from failing, make sure that the following conditions apply:

- The disk has two free partitions for the public and private regions.
- The disk has an `s2` slice.
- The disk has a small amount of free space (at least 1 megabyte at the beginning or end of the disk) that does not belong to any partition. If the disk being encapsulated is the root disk, and this does not have sufficient free space available, a similar sized portion of the swap partition is used instead.

Only encapsulate a root disk if you also intend to mirror it. There is no benefit in root-disk encapsulation for its own sake.

See [“Rootability”](#) on page 872.

Use the `format` or `fdisk` commands to obtain a printout of the root disk partition table before you encapsulate a root disk. For more information, see the appropriate manual pages. You may need this information should you subsequently need to recreate the original root disk.

You cannot grow or shrink any volume (`rootvol`, `usrvol`, `varvol`, `optvol`, `swapvol`, and so on) that is associated with an encapsulated root disk. This is because these volumes map to physical partitions on the disk, and these partitions must be contiguous.

When the boot disk is encapsulated or mirrored, a device path alias is added to the NVRAMRC in the SPARC EEPROM. These device aliases can be used to set the system's boot device.

For more information, see the `devalias` and `boot-device` settings in the SUN documentation.

Warning: If the root disk is encapsulated and the dump device is covered by the swap volume, it is not safe to use the `savecore -L` operation because this overwrites the swap area. Configure a dedicated dump device on a partition other than the swap area.

To encapsulate a disk for use in VxVM

- 1 Select `Encapsulate` one or more disks from the `vxdiskadm` main menu.

Your system may use device names that differ from the examples shown here.

At the following prompt, enter the disk device name for the disks to be encapsulated:

```
Select disk devices to encapsulate:
[<pattern-list>,all,list,q,?] device name
```

The *pattern-list* can be a single disk, or a series of disks and/or controllers (with optional targets). If *pattern-list* consists of multiple items, those items must be separated by white space.

If you do not know the address (device name) of the disk to be encapsulated, enter `l` or `list` at the prompt for a complete listing of available disks.

- 2 To continue the operation, enter `y` (or press Return) at the following prompt:

```
Here is the disk selected. Output format: [Device]
device name
```

```
Continue operation? [y,n,q,?] (default: y) y
```

- 3 Select the disk group to which the disk is to be added at the following prompt:

You can choose to add this disk to an existing disk group or to a new disk group. To create a new disk group, select a disk group name that does not yet exist.

```
Which disk group [<group>,list,q,?]
```

- 4 At the following prompt, either press Return to accept the default disk name or enter a disk name:

```
Use a default disk name for the disk? [y,n,q,?] (default: y)
```

- 5 To continue with the operation, enter `y` (or press Return) at the following prompt:

The selected disks will be encapsulated and added to the *disk group name* disk group with default disk names.

```
device name
```

```
Continue with operation? [y,n,q,?] (default: y) y
```


- 6 To confirm that encapsulation should proceed, enter `y` (or press Return) at the following prompt:

```
The following disk has been selected for encapsulation.
Output format: [Device]
```

```
device name
```

```
Continue with encapsulation? [y,n,q,?] (default: y) y
```

A message similar to the following confirms that the disk is being encapsulated for use in VxVM and tells you that a reboot is needed:

```
The disk device device name will be encapsulated and added to
the disk group diskgroup with the disk name diskgroup01.
```

- 7 For non-root disks, you can now choose whether the disk is to be formatted as a CDS disk that is portable between different operating systems, or as a non-portable sliced disk:

```
Enter the desired format [cdsdisk,sliced,q,?]
(default: cdsdisk)
```

Enter the format that is appropriate for your needs. In most cases, this is the default format, `cdsdisk`. Note that only the `sliced` format is suitable for use with root, boot or swap disks.

- 8 At the following prompt, `vxdiskadm` asks if you want to use the default private region size of 65536 blocks (32MB). Press Return to confirm that you want to use the default value, or enter a different value. (The maximum value that you can specify is 524288 blocks.)

```
Enter desired private region length [<privlen>,q,?]
(default: 65536)
```

- 9 If you entered `cdsdisk` as the format in step 7, you are prompted for the action to be taken if the disk cannot be converted this format:

```
Do you want to use sliced as the format should cdsdisk
fail? [y,n,q,?] (default: y)
```

If you enter `y`, and it is not possible to encapsulate the disk as a CDS disk, it is encapsulated as a sliced disk. Otherwise, the encapsulation fails.

- 10 `vxdiskadm` then proceeds to encapsulate the disks. You should now reboot your system at the earliest possible opportunity, for example by running this command:

```
# shutdown -g0 -y -i6
```

The `/etc/vfstab` file is updated to include the volume devices that are used to mount any encapsulated file systems. You may need to update any other references in backup scripts, databases, or manually created swap devices. The original `/etc/vfstab` file is saved as `/etc/vfstab.prevm`.

- 11 At the following prompt, indicate whether you want to encapsulate more disks (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Encapsulate other disks? [y,n,q,?] (default: n) n
```

The default layout that is used to encapsulate disks can be changed.

Failure of disk encapsulation

Under some circumstances, encapsulation of a disk can fail because there is not enough free space available on the disk to accommodate the private region. If there is insufficient free space, the encapsulation process ends abruptly with an error message similar to the following:

```
VxVM ERROR V-5-2-338 The encapsulation operation failed with the
following error:
It is not possible to encapsulate device, for the following
reason:
<VxVM vxslicer ERROR V-5-1-1108 Unsupported disk layout.>
```

One solution is to configure the disk with the `nopriv` format.

See [“Using nopriv disks for encapsulation”](#) on page 870.

Using nopriv disks for encapsulation

Encapsulation converts existing partitions on a specified disk to volumes. If any partitions contain file systems, their `/etc/vfstab` entries are modified so the file systems are mounted on volumes instead.

Disk encapsulation requires that enough free space be available on the disk (by default, 32 megabytes) for storing the private region that VxVM uses for disk identification and configuration information. This free space cannot be included in any other partitions.

See the `vxencap(1M)` manual page.

You can encapsulate a disk that does not have space available for the VxVM private region partition by using the `vxdisk` utility. To do this, configure the disk as a `nopriv` device that does not have a private region.

The drawback with using `nopriv` devices is that VxVM cannot track changes in the address or controller of the disk. Normally, VxVM uses identifying information stored in the private region on the physical disk to track changes in the location of a physical disk. Because `nopriv` devices do not have private regions and have no identifying information stored on the physical disk, tracking cannot occur.

One use of `nopriv` devices is to encapsulate a disk so that you can use VxVM to move data off the disk. When space has been made available on the disk, remove the `nopriv` device, and encapsulate the disk as a standard disk device.

A disk group cannot be formed entirely from `nopriv` devices. This is because `nopriv` devices do not provide space for storing disk group configuration information. Configuration information must be stored on at least one disk in the disk group.

Creating a nopriv disk for encapsulation

Warning: Do not use `nopriv` disks to encapsulate a root disk. If insufficient free space exists on the root disk for the private region, part of the swap area can be used instead.

To create a nopriv disk for encapsulation

- 1 If it does not exist already, set up a partition on the disk for the area that you want to access using VxVM.
- 2 Use the following command to map a VM disk to the partition:

```
# vxdisk define partition-device type=nopriv
```

where *partition-device* is the basename of the device in the `/dev/dsk` directory.

For example, to map partition 3 of disk device `c0t4d0`, use the following command:

```
# vxdisk define c0t4d0s3 type=nopriv
```

Creating volumes for other partitions on a nopriv disk

To create volumes for other partitions on a nopriv disk

- 1 Add the partition to a disk group.
- 2 Determine where the partition resides within the encapsulated partition.
- 3 If no data is to be preserved on the partition, use `vxassist` to create a volume with the required length.

Warning: By default, `vxassist` re-initializes the data area of a volume that it creates. If there is data to be preserved on the partition, do not use `vxassist`. Instead, create the volume with `vxmake` and start the volume with the command `vxvol init active`.

Rootability

VxVM can place various files from the root file system, `swap` device, and other file systems on the root disk under VxVM control. This is called rootability. The root disk (that is, the disk containing the root file system) can be put under VxVM control through the process of encapsulation.

The root disk can be encapsulated using the `vxdiskadm` command.

See [“Encapsulating a disk”](#) on page 866.

Once encapsulated, the root disk can also be mirrored by using the `vxdiskadm` command.

See [“Mirroring an encapsulated root disk”](#) on page 875.

Encapsulation converts existing partitions on that disk to volumes. Once under VxVM control, the `root` and `swap` devices appear as volumes and provide the same characteristics as other VxVM volumes. A volume that is configured for use as a swap area is referred to as a swap volume, and a volume that contains the root file system is referred to as a root volume.

Note: Only encapsulate your root disk if you also intend to mirror it. There is no benefit in root-disk encapsulation for its own sake.

You can mirror the `rootvol`, and `swapvol` volumes, as well as other parts of the root disk that are required for a successful boot of the system (for example, `/usr`). This provides complete redundancy and recovery capability in the event of disk

failure. Without mirroring, the loss of the `root`, `swap`, or `usr` partition prevents the system from being booted from surviving disks.

Mirroring disk drives that are critical to booting ensures that no single disk failure renders the system unusable. A suggested configuration is to mirror the critical disk onto another available disk (using the `vxdiskadm` command). If the disk containing `root` and `swap` partitions fails, the system can be rebooted from a disk containing mirrors of these partitions.

Recovering a system after the failure of an encapsulated root disk requires the application of special procedures.

See the *Veritas Volume Manager Troubleshooting Guide*.

Booting root volumes

When the operating system is booted, the `root` file system and `swap` area must be available for use before the `vxconfigd` daemon can load the VxVM configuration or start any volumes. During system startup, the operating system must see the `rootvol` and `swapvol` volumes as regular partitions so that it can access them as ordinary disk partitions.

Due to this restriction, each of the `rootvol` and `swapvol` plexes must be created from contiguous space on a disk that is mapped to a single partition. It is not possible to stripe, concatenate or span the plex of a `rootvol` or `swapvol` volume that is used for booting. Any mirrors of these plexes that are potentially bootable also cannot be striped, concatenated or spanned.

For information on how to configure your system BIOS to boot from a disk other than the default boot disk, refer to the documentation from your hardware vendor.

Boot-time volume restrictions

When the operating system is booted, the `root` file system and `swap` area must be available for use before the `vxconfigd` daemon can load the VxVM configuration or start any volumes. During system startup, the operating system must see the `rootvol` and `swapvol` volumes as regular partitions so that it can access them as ordinary disk partitions.

Due to this restriction, each of the `rootvol` and `swapvol` plexes must be created from contiguous space on a disk that is mapped to a single partition. It is not possible to stripe, concatenate or span the plex of a `rootvol` or `swapvol` volume that is used for booting. Any mirrors of these plexes that are potentially bootable also cannot be striped, concatenated or spanned.

Volumes on the root disk have the following restrictions on their configuration:

- For the x64 platform, root encapsulation is supported for Update 1 and later releases of the Solaris 10 OS that include the GRUB boot loader.
- The root volume (`rootvol`) must exist in the disk group that is chosen to be the boot disk group, `bootdg`. Although other volumes named `rootvol` can be created in other disk groups, only the `rootvol` in `bootdg` can be used to boot the system.
- If the volumes that are required to boot the system span multiple disks, all these disks must be in the boot disk group.
- The `rootvol` and `swapvol` volumes always have minor device numbers 0 and 1 respectively. Other volumes on the root disk do not have specific minor device numbers.
- Restricted mirrors of volumes on the root disk have overlay partitions created for them. An overlay partition is one that exactly includes the disk space occupied by the restricted mirror. During boot, before the `rootvol`, `varvol`, `usrvol` and `swapvol` volumes are fully configured, the default volume configuration uses the overlay partition to access the data on the disk.
- Although it is possible to add a striped mirror to a `rootvol` device for performance reasons, you cannot stripe the primary plex or any mirrors of `rootvol` that may be needed for system recovery or booting purposes if the primary plex fails.
- `rootvol` and `swapvol` cannot be spanned or contain a primary plex with multiple noncontiguous subdisks. You cannot grow or shrink any volume associated with an encapsulated boot disk (`rootvol`, `usrvol`, `varvol`, `optvol`, `swapvol`, and so on) because these map to a physical underlying partition on the disk and must be contiguous. A workaround is to unencapsulate the boot disk, repartition the boot disk as desired (growing or shrinking partitions as needed), and then re-encapsulating.
- When mirroring parts of the boot disk, the disk being mirrored to must be large enough to hold the data on the original plex, or mirroring may not work.
- The volumes on the root disk cannot use dirty region logging (DRL).

In addition to these requirements, it is a good idea to have at least one contiguous, (cylinder-aligned if appropriate) mirror for each of the volumes for `root`, `usr`, `var`, `opt` and `swap`. This makes it easier to convert these from volumes back to regular disk partitions (during an operating system upgrade, for example).

Creating redundancy for the root disk

You can create an active backup of the root disk, in case of a single disk failure. Use the `vxrootadm` command to create a mirror of the booted root disk, and other volumes in the root disk group.

To create a back-up root disk

- ◆ Create a mirror with the `vxrootadm addmirror` command.

```
# vxrootadm [-v] [-Y] addmirror targetdisk
```

Creating an archived back-up root disk for disaster recovery

In addition to having an active backup of the root disk, you can keep an archived back-up copy of the bootable root disk. Use the `vxrootadm` command to create a snapshot of the booted root disk, which creates a mirror and breaks it off into a separate disk group.

To create an archived back-up root disk

- 1 Add a disk to the booted root disk group.
- 2 Create a snapshot of the booted root disk.

```
# vxrootadm [-v] mksnap targetdisk targetdg
```

- 3 Archive the back-up root disk group for disaster recovery.

Mirroring an encapsulated root disk

VxVM allows you to mirror the root volume and other areas needed for booting onto another disk. Mirroring the root volume enables you to recover from a failure of your root disk by replacing it with one of its mirrors.

For Sun x64 systems, mirroring a root disk creates a GRUB boot menu entry for the Primary and Alternate (mirror) Boot disk.

For Sun SPARC systems, after mirroring the root disk, you can configure the system to boot from the alternate boot drive to recover from a primary boot drive failure.

See the *Storage Foundation High Availability Solutions Troubleshooting Guide* for more information about recovering from boot drive failure.

To mirror your root disk onto another disk

- 1 Choose a disk that is at least as large as the existing `root` disk.
- 2 If the selected disk is not already under VxVM control, use the `vxdiskadd` or the `vxdiskadm` command to add it to the `bootdg` disk group. Ensure that you specify the `sliced` format for the disk.

- 3 Select **Mirror Volumes on a Disk** from the `vxdiskadm` main menu to create a mirror of the root disk. Doing so automatically invokes the `vxmirror` command if the mirroring operation is performed on the `root` disk.

Alternatively, to mirror only those file systems on the root disk that are required to boot the system, run the following command:

```
# vxmirror boot_disk altboot_disk
```

where `altboot_disk` is the disk media name of the mirror for the root disk.

`vxmirror` creates a mirror for `rootvol` (the volume for the `root` file system on an alternate disk). The alternate `root` disk is configured to enable booting from it if the primary `root` disk fails.

- 4 Monitor the progress of the mirroring operation with the `vxtask list` command.

```
# vxtask list
TASKID PTID TYPE/STATE PCT PROGRESS
161 PARENT/R 0.00% 3/0(1) VXRECOVER dg01 dg
162 162 ATCOPY/R 04.77% 0/41945715/2000896 PLXATT home home-01 dg
```

Booting from alternate boot disks

If the root disk is encapsulated and mirrored, you can use one of its mirrors to boot the system if the primary boot disk fails. This procedure differs between Solaris SPARC systems and Solaris x64 systems.

Booting from an alternate (mirror) boot disk on Solaris x64 systems

On a Solaris x64 system, the alternate boot disk is added to the GRUB boot menu when a boot disk is mirrored. If one root disk fails, the system stays up and lets you replace the disk. No reboot is required to perform this maintenance with internal SAS controllers and other CRU-type drives that are hot swappable. Replace the disk, then rescan with the `vxctl enable` command to discover the replacement.

Alternatively, the `bootpath` can be redefined in the EEPROM without changing the GRUB configuration.

See [“The boot process on x64 systems”](#) on page 877.

Console access and the ability to select from the GRUB menu is required for the following procedure. The system should not have rebooted because of plex failure, but may have rebooted for other reasons.

To boot from an alternate boot disk on a Solaris x64 system

1 Select the "Alternate" GRUB menu entry:

```
title Solaris 10 11/06 s10x_u3wos_10 x64 <VxVM: Alternate Boot Disk>
    root (hd0,0,a)
    kernel /platform/i64pc/multiboot
    module /platform/i64pc/boot_archive.alt
```

2 After the system has booted, replace the failed drive.

See the *Veritas Storage Foundation and High Availability Solutions Troubleshooting Guide*.

The boot process on x64 systems

From Update 1 of the Solaris 10 OS, x64 systems are configured to use the GRUB boot loader. The devices from which a system may be booted are defined in the GRUB configuration file, `/boot/grub/menu.lst`. From the GRUB menu, you can select from the available bootable partitions that are known to the system. By default, the system will boot from the device that is defined by the `bootpath` variable in the EEPROM.

Defining root disk mirrors as bootable

After creating a root disk mirror, you can make it available for booting.

On Sun x64 systems, VxVM automatically creates a GRUB menu entry for the alternate boot disk when the boot disk is mirrored. During the booting process, select the alternate GRUB menu entry from the system console.

An alternate method is to change the 'default' GRUB menu setting in the `/boot/grub/menu.lst` file to select this entry automatically during the booting process.

For details, see the *Veritas Storage Foundation and High Availability Troubleshooting Guide*.

Booting from an alternate boot disk on Solaris SPARC systems

If the root disk is encapsulated and mirrored, you can use one of its mirrors to boot the system if the primary boot disk fails.

On a Solaris SPARC system, booting from an alternate boot disk requires that some EEPROM settings are changed.

See [“The boot process on Solaris SPARC systems”](#) on page 878.

The boot process on Solaris SPARC systems

A Solaris SPARC® system prompts for a boot command unless the `autoboot` flag has been set in the nonvolatile storage area used by the firmware. Machines with older PROMs have different prompts than that for the newer V2 and V3 versions. These newer versions of PROM are also known as OpenBoot PROMs (OBP). The `boot` command syntax for the newer types of PROMs is:

```
ok boot [OBP names] [filename] [boot-flags]
```

The OBP names specify the OpenBoot PROM designations. For example, on Desktop SPARC systems, the designation `sbus/esp@0,800000/sd@3,0:a` indicates a SCSI disk (`sd`) at target 3, lun 0 on the SCSI bus, with the `esp` host bus adapter plugged into slot 0.

You can use Veritas Volume Manager boot disk alias names instead of OBP names. Example aliases are `vx-rootdisk` or `vx-disk01`. To list the available boot devices, use the `devalias` command at the OpenBoot prompt.

The filename argument is the name of a file that contains the kernel. The default is `/kernel/unix` in the `root` partition. If necessary, you can specify another program (such as `/stand/diag`) by specifying the `-a` flag. (Some versions of the firmware allow the default filename to be saved in the nonvolatile storage area of the system.)

Warning: Do not boot a system running VxVM with rootability enabled using all the defaults presented by the `-a` flag.

Boot flags are not interpreted by the `boot` program. The `boot` program passes all boot-flags to the file identified by filename.

See the `kadb` (1M) manual page.

See the `kernel` (1) manual page.

Mirroring other file systems on the root disk

There may be other volumes on the root disk, such as volumes for `/home` or `/tmp` file systems. If necessary, these can be mirrored separately using the `vxassist` utility. For example, if you have a `/home` file system on a volume `homevol`, you can mirror it to `alternate_disk` using the command:

```
# vxassist mirror homevol alternate_disk
```

If you do not have space for a copy of some of these file systems on your alternate boot disk, you can mirror them to other disks. You can also span or stripe these other volumes across other disks attached to your system.

To list all volumes on your primary boot disk, use the command:

```
# vxprint -t -v -e 'aslist.aslist.sd_disk="boot_disk"'
```

Encapsulating SAN disks

Most fourth generation servers do not have local storage. They support booting from a SAN; that is, booting directly from a multi-pathed storage array.

A Solaris system may be booted from a SAN disk under the following conditions:

- For Solaris 10, the operating system can be installed directly onto a fabric disk in a SAN environment.

Veritas Volume Manager can encapsulate a bootable SAN disk provided that the disk is listed as being supported for this purpose in the Hardware Compatibility List (HCL):

<http://www.symantec.com/docs/TECH170013>

For some disk arrays, special hardware configuration may be required to allow a system to be booted from one of the LUNs in the array. Refer to the documentation supplied by the array vendor for more information. Having configured the disk array so that you can boot your system from it, you can proceed to encapsulate it using VxVM.

To migrate from an internal boot disk to a SAN boot disk:

- 1 Verify that the HCL lists the target SAN disk as being supported for SAN booting.
- 2 Use `Add or initialize one or more disks` from the `vxdiskadm` main menu to add the target SAN disk to the boot disk group (aliased as `bootdg`).
- 3 Use `Mirror Volumes on a Disk` from the `vxdiskadm` main menu to create a mirror of the root disk on the target disk.
- 4 Boot from the mirror disk to verify that the system is still bootable.

Once you have booted the system from the SAN disk, you can mirror it to another SAN disk that has been added to the boot disk group.

If required, you can remove the plexes of the original boot disk by using the `vxplex` command. For example, the following command removes the plexes `rootvol-01`, `swapvol-01`, and `home-01` that are configured on the boot disk:

```
# vxplex -o rm dis rootvol-01 swapvol-01 home-01
```

Best practices and guidelines for booting from SAN LUNS

When using an A/A array to boot from a SAN, complete the following steps in the order presented:

- Configure SAN boot LUNs with the lowest host target id. For example, target id 0, 1.
- Encapsulate the root disk (SAN LUN) into VxVM.
- Mirror the encapsulated root disk to another LUN on the same or different enclosure.
- Create boot aliases/sequence so that server will boot from an alternate path if the primary boot path is not available.

When using A/P, A/P-F, and ALUA arrays to boot from a SAN, the following additional best practices need to be applied.

- If the SAN Boot LUN and its mirror are on the same enclosure, they should have the same default owner (A/P array) or optimized path (ALUA) on the same controller. Otherwise, there are issues with "mirroring" the SAN boot LUN.
- Boot aliases/sequence should be configured on the paths of the array controller which is the LUN owner (A/P and A/PF) and optimized path (ALUA) for the SAN boot LUNs and its mirror.

Administering an encapsulated boot disk

The `vxrootadm` command lets you make a bootable snapshot of an encapsulated boot disk or grow the encapsulated boot disk.

The `vxrootadm` command has the following format:

```
# vxrootadm [-v] [-g dg] [-s srcdisk] ... keyword arg ...
```

Valid keyword and argument combinations are as follows:

- # `vxrootadm -s srcdisk mksnap destdisk newdg`
- # `vxrootadm -s srcdisk grow destdisk volumename=newsize ...`
- # `vxrootadm grow continue`

See [“Growing an encapsulated boot disk”](#) on page 881.

`vxrootadm` includes the following options:

```
vxrootadm [-v ] [-D]
```

These are verbose and debug message options and are optional.

```
vxrootadm [-g dg]
```

The disk group argument is optional; however, it is only used with the `mksnap` keyword or during phase 1 of a grow operation.

```
vxrootadm -s srcdisk
```

Specifies the source disk.

See the `vxrootadm(1m)` manual page.

Creating a snapshot of an encapsulated boot disk

To create a snapshot of an encapsulated boot disk, use the `vxrootadm` command. The target disk for the snapshot must be as large (or bigger) than the source disk (boot disk). You must use a new disk group name to associate the target disk.

To create a snapshot of an encapsulated boot disk

- ◆ Enter the following command:

```
# vxrootadm -s srcdisk [-g dg] mksnap destdisk newdg
```

For example:

```
# vxrootadm -s disk_0 -g rootdg mksnap disk_1 snapdg
```

In this example, `disk_0` is the encapsulated boot disk, and `rootdg` is the associate boot disk group. `disk_1` is the target disk, and `snapdg` is the new disk group name

See [“Booting from alternate boot disks”](#) on page 876.

Growing an encapsulated boot disk

The target disk for the grow operation must be large enough to accommodate the grown volumes. The grow operation can be performed on the active boot disk or a snapshot boot disk.

When you grow an encapsulated boot disk, the format of the `grow` keyword depends on whether the operation can be completed in one or more phases.

To grow a non-booted root disk created as a snapshot or using the `mksnap` keyword, use the following format.

```
# vxrootadm -s srcdisk [-g dg] grow destdisk volumename=newsize \
volumename=newsize ...
```

Because the disk is not the currently booted root disk, you can complete all operations in a single phase without a reboot.

Growing a booted root disk requires four phases. For phase 1, use the command above. For phases 2 to 4, specify `vxrootadm grow continue`.

To grow an active encapsulated boot disk

- 1 To complete the grow operation on the active boot disk requires three reboots to complete the grow operations for the selected volume (`rootvol`, `usrvol`, or `swapvol`).
- 2 Enter the following command:

```
# vxrootadm -s disk_0 -g rootdg grow disk_1 rootvol=80g
```

In this example, `disk_0` is an encapsulated boot disk associated with the boot disk group `rootdg`. `disk_1` is the target disk and `rootvol` is a 60g volume to be grown.

You are prompted when a reboot is required (with specific command needed), and how to continue the grow operation after the reboot is completed.

When the grow operation completes, the target disk is the active boot disk, the volume has grown to the selected size, and the source boot disk is removed from the boot disk group (`rootdg`).

To grow a snapshot (not active) encapsulated boot disk

- 1 For a grow operation on the snapshot (not active) boot disk, no reboots are required to complete the grow operations for the selected volume (`rootvol`, `usrvol`, or `swapvol`).
- 2 Enter the following command:

```
# vxrootadm -s disk_1 -g snapdg grow disk_2 rootvol=80g
```

In this example, `disk_1` is a snapshot encapsulated boot disk associated with the disk group `snapdg`. `disk_2` is the target disk and `rootvol` is a 60g volume to be grown.

When the grow operation completes, the target disk volume is grown to the indicated size, and the source disk is removed from the disk group (`snapdg`).

See [“Booting from alternate boot disks”](#) on page 876.

Unencapsulating the root disk

You can use the `vxunroot` utility to remove rootability support from a system. This makes `root`, `swap`, `home` and other file systems on the root disk directly accessible through disk partitions, instead of through volume devices.

The `vxunroot` utility also makes the necessary configuration changes to allow the system to boot without any dependency on VxVM.

Only the volumes that were present on the root disk when it was encapsulated can be unencapsulated using `vxunroot`. Before running `vxunroot`, evacuate all other volumes that were created on the root disk after it was encapsulated.

Do not remove the plexes on the root disk that correspond to the original disk partitions.

Warning: This procedure requires a reboot of the system.

To remove rootability from a system

- 1 Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` on the disks other than the root disk.

For example, the following command removes the plexes `rootvol-02`, `swapvol-02`, and `home-02` that are configured on the boot disk mirror:

```
# vxplex -g bootdg -o rm dis rootvol-02 swapvol-02 home-02
```

- 2 Run the `vxunroot` utility:

```
# vxunroot
```

`vxunroot` does not perform any conversion to disk partitions if any plexes remain on other disks.

Managing tunable parameters

This chapter includes the following topics:

- [Reference tables of Storage Foundation tunable parameters](#)
- [How to change tunable parameters](#)

Reference tables of Storage Foundation tunable parameters

Refer to the table for the Storage Foundation component that you want to tune.

DMP tunable parameters

DMP provides various parameters that you can use to tune your environment.

[Table 39-1](#) shows the DMP parameters that can be tuned. You can set a tunable parameter online, without a reboot.

Table 39-1 DMP parameters that are tunable

Parameter	Description
dmp_cache_open	If this parameter is set to <code>on</code>, the first open of a device that is performed by an array support library (ASL) is cached. This caching enhances the performance of device discovery by minimizing the overhead that is caused by subsequent opens by ASLs. If this parameter is set to <code>off</code>, caching is not performed. The default value is <code>on</code>. DMP support for native multi-pathing requires that the caching is <code>on</code>. Therefore, you cannot set the <code>dmp_cache_open</code> parameter to <code>off</code> if the <code>dmp_native_multipathing</code> parameter is set to <code>on</code>.
dmp_daemon_count	The number of kernel threads that are available for servicing path error handling, path restoration, and other DMP administrative tasks. The default number of threads is 10.
dmp_delayq_interval	How long DMP should wait before retrying I/O after an array fails over to a standby path. Some disk arrays are not capable of accepting I/O requests immediately after failover. The default value is 15 seconds.
dmp_fast_recovery	Whether DMP should try to obtain SCSI error information directly from the HBA interface. Setting the value to <code>on</code> can potentially provide faster error recovery, if the HBA interface supports the error enquiry feature. If this parameter is set to <code>off</code>, the HBA interface is not used. The default setting is <code>on</code>.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_health_time</code>	DMP detects intermittently failing paths, and prevents I/O requests from being sent on them. The value of <code>dmp_health_time</code> represents the time in seconds for which a path must stay healthy. If a path's state changes back from enabled to disabled within this time period, DMP marks the path as intermittently failing, and does not re-enable the path for I/O until <code>dmp_path_age</code> seconds elapse. The default value is 60 seconds. A value of 0 prevents DMP from detecting intermittently failing paths.
<code>dmp_log_level</code>	The level of detail that is displayed for DMP console messages. The following level values are defined: 1 — Displays all DMP log messages that existed in releases before 5.0. 2 — Displays level 1 messages plus messages that relate to path or disk addition or removal, SCSI errors, IO errors and DMP node migration. 3 — Displays level 1 and 2 messages plus messages that relate to path throttling, suspect path, idle path and insane path logic. 4 — Displays level 1, 2 and 3 messages plus messages that relate to setting or changing attributes on a path and tunable related changes. The default value is 1.
<code>dmp_low_impact_probe</code>	Determines if the path probing by restore daemon is optimized or not. Set it to <code>on</code> to enable optimization and <code>off</code> to disable. Path probing is optimized only when restore policy is <code>check_disabled</code> or during <code>check_disabled</code> phase of <code>check_periodic</code> policy. The default value is <code>on</code>.

Table 39-1 DMP parameters that are tunable (continued)

Parameter	Description
dmp_lun_retry_timeout	Specifies a retry period for handling transient errors that are not handled by the HBA and the SCSI driver. In general, no such special handling is required. Therefore, the default value of the <code>dmp_lun_retry_timeout</code> tunable parameter is 0. When all paths to a disk fail, DMP fails the I/Os to the application. The paths are checked for connectivity only once. In special cases when DMP needs to handle the transient errors, configure DMP to delay failing the I/Os to the application for a short interval. Set the <code>dmp_lun_retry_timeout</code> tunable parameter to a non-zero value to specify the interval. If all of the paths to the LUN fail and I/Os need to be serviced, then DMP probes the paths every five seconds for the specified interval. If the paths are restored within the interval, DMP detects this and retries the I/Os. DMP does not fail I/Os to a disk with all failed paths until the specified <code>dmp_lun_retry_timeout</code> interval or until the I/O succeeds on one of the paths, whichever happens first.
dmp_monitor_fabric	Determines whether the Event Source daemon (<code>vxesd</code>) uses the Storage Networking Industry Association (SNIA) HBA API. This API allows DDL to improve the performance of failover by collecting information about the SAN topology and by monitoring fabric events. If this parameter is set to <code>on</code>, DDL uses the SNIA HBA API. Note that the HBA vendor specific HBA-API library should be available to use this feature. If this parameter is set to <code>off</code>, the SNIA HBA API is not used. The default setting is <code>off</code> for releases before 5.0 that have been patched to support this DDL feature. The default setting is <code>on</code> for 5.0 and later releases.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_monitor_osevent</code>	Determines whether the Event Source daemon (<code>vxesd</code>) monitors operating system events such as reconfiguration operations. If this parameter is set to <code>on</code>, <code>vxesd</code> monitors operations such as attaching operating system devices. If this parameter is set to <code>off</code>, <code>vxesd</code> does not monitor operating system operations. When DMP co-exists with EMC PowerPath, Symantec recommends setting this parameter to <code>off</code> to avoid any issues. The default setting is <code>on</code>, unless EMC PowerPath is installed. If you install DMP on a system that already has PowerPath installed, DMP sets the <code>dmp_monitor_osevent</code> to <code>off</code>.
<code>dmp_monitor_ownership</code>	Determines whether the ownership monitoring is enabled for ALUA arrays. When this tunable is set to <code>on</code>, DMP polls the devices for LUN ownership changes. The polling interval is specified by the <code>dmp_restore_interval</code> tunable. The default value is <code>on</code>. When the <code>dmp_monitor_ownership</code> tunable is <code>off</code>, DMP does not poll the devices for LUN ownership changes.
<code>dmp_native_multipathing</code>	Determines whether DMP will intercept the I/Os directly on the raw OS paths or not. Set the tunable to <code>on</code> to have DMP do multi-pathing of I/Os done directly on raw paths, otherwise set it to <code>off</code>. The default value is <code>off</code>. DMP support for native multi-pathing requires that the caching is <code>on</code>. Therefore, you cannot set the <code>dmp_native_multipathing</code> to <code>on</code> if the <code>dmp_cache_open</code> parameter is set to <code>off</code>.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_native_support</code>	Determines whether DMP will do multi-pathing for native devices. Set the tunable to <code>on</code> to have DMP do multi-pathing for native devices. When a Storage Foundation Cluster File System High Availability product is installed, the default value is <code>off</code>. When Veritas Dynamic Multi-Pathing is installed, the default value is <code>on</code>.
<code>dmp_path_age</code>	The time for which an intermittently failing path needs to be monitored as healthy before DMP again tries to schedule I/O requests on it. The default value is 300 seconds. A value of 0 prevents DMP from detecting intermittently failing paths.
<code>dmp_pathswitch_blks_shift</code>	The default number of contiguous I/O blocks that are sent along a DMP path to an array before switching to the next available path. The value is expressed as the integer exponent of a power of 2; for example 9 represents 512 blocks. The default value is 9. In this case, 512 blocks (256k) of contiguous I/O are sent over a DMP path before switching. For intelligent disk arrays with internal data caches, better throughput may be obtained by increasing the value of this tunable. For example, for the HDS 9960 A/A array, the optimal value is between 15 and 17 for an I/O activity pattern that consists mostly of sequential reads or writes. This parameter only affects the behavior of the <code>balanced</code> I/O policy. A value of 0 disables multi-pathing for the policy unless the <code>vxddmpadm</code> command is used to specify a different partition size for an array. See “Specifying the DMP I/O policy” on page 292.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_probe_idle_lun</code>	If DMP statistics gathering is enabled, set this tunable to <code>on</code> (default) to have the DMP path restoration thread probe idle LUNs. Set this tunable to <code>off</code> to turn off this feature. (Idle LUNs are VM disks on which no I/O requests are scheduled.) The value of this tunable is only interpreted when DMP statistics gathering is enabled. Turning off statistics gathering also disables idle LUN probing. The default value is <code>on</code>.
<code>dmp_probe_threshold</code>	If the <code>dmp_low_impact_probe</code> is turned <code>on</code>, <code>dmp_probe_threshold</code> determines the number of paths to probe before deciding on changing the state of other paths in the same subpath failover group. The default value is 5.
<code>dmp_restore_cycles</code>	If the DMP restore policy is <code>check_periodic</code>, the number of cycles after which the <code>check_all</code> policy is called. The default value is 10. The value of this tunable can also be set using the <code>vxdmpadm start restore</code> command. See “Configuring DMP path restoration policies” on page 305.
<code>dmp_restore_interval</code>	The interval attribute specifies how often the path restoration thread examines the paths. Specify the time in seconds. The default value is 300. The value of this tunable can also be set using the <code>vxdmpadm start restore</code> command. See “Configuring DMP path restoration policies” on page 305.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
dmp_restore_policy	The DMP restore policy, which can be set to one of the following values: ■ <code>check_all</code> ■ <code>check_alterate</code> ■ <code>check_disabled</code> ■ <code>check_periodic</code> The default value is <code>check_disabled</code>. The value of this tunable can also be set using the <code>vxddmpadm start restore</code> command. See “Configuring DMP path restoration policies” on page 305.
dmp_restore_state	If this parameter is set to <code>enabled</code>, it enables the path restoration thread to be started. See “Configuring DMP path restoration policies” on page 305. If this parameter is set to <code>disabled</code>, it stops and disables the path restoration thread. If this parameter is set to <code>stopped</code>, it stops the path restoration thread until the next device discovery cycle. The default is <code>enabled</code>. See “Stopping the DMP path restoration thread” on page 307.
dmp_retry_count	When I/O fails on a path with a path busy error, DMP marks the path as busy and avoids using it for the next 15 seconds. If a path reports a path busy error for <code>dmp_retry_count</code> number of times consecutively, DMP marks the path as failed. The default value of <code>dmp_retry_count</code> is 5.
dmp_scsi_timeout	Determines the timeout value to be set for any SCSI command that is sent via DMP. If the HBA does not receive a response for a SCSI command that it has sent to the device within the timeout period, the SCSI command is returned with a failure error code. The default value is 30 seconds.

Table 39-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_sfg_threshold</code>	Determines the minimum number of paths that should be failed in a failover group before DMP starts suspecting other paths in the same failover group. The value of 0 disables the failover logic based on subpath failover groups. The default value is 1.
<code>dmp_stat_interval</code>	The time interval between gathering DMP statistics. The default and minimum value are 1 second.

Tunable parameters for VxVM

Veritas Volume Manager (VxVM) has several parameters that you can use to tune the environment. The VxVM tunable parameters comprise several components.

The VxVM components for tunable parameters are as follows:

- **basevm**
Parameters to tune the core functionality of VxVM.
See [“Tunable parameters for core VxVM”](#) on page 892.
- **cvm**
Parameters to tune Cluster Volume Manager (CVM).
See [“Tunable parameters for CVM”](#) on page 904.
- **fmr**
Parameters to tune the FlashSnap functionality (FMR).
See [“Tunable parameters for FlashSnap \(FMR\)”](#) on page 898.
- **vvr**
Parameters to tune Veritas Volume Replicator (VVR).
See [“Tunable parameters for VVR”](#) on page 904.

Tunable parameters for core VxVM

[Table 39-2](#) lists the kernel tunable parameters for VxVM.

You can tune the parameters using the `vxtune` command or the operating system method, unless otherwise noted.

Table 39-2 Kernel tunable parameters for core VxVM

Parameter	Description
<code>vol_checkpoint_default</code>	The interval at which utilities performing recoveries or resynchronization operations load the current offset into the kernel as a checkpoint. A system failure during such operations does not require a full recovery, but can continue from the last reached checkpoint. The default value is 20480 sectors (10MB). Increasing this size reduces the overhead of checkpoints on recovery operations at the expense of additional recovery following a system failure during a recovery.
<code>vol_default_iodelay</code>	The count in clock ticks for which utilities pause if they have been directed to reduce the frequency of issuing I/O requests, but have not been given a specific delay time. This tunable is used by utilities performing operations such as resynchronizing mirrors or rebuilding RAID-5 columns. The default value is 50 ticks. Increasing this value results in slower recovery operations and consequently lower system impact while recoveries are being performed.
<code>vol_kmsg_resend_period</code>	This parameter cannot be tuned with the <code>vxtune</code> command. If specified in the <code>/kernel/drv/vxio.conf</code> file, the value is internally converted to microseconds. The value is then applied to the <code>vol_kmsg_resend_period_usecs</code> parameter, which represents the value of the kernel message (KMSG) resend period that is used by the clustering functionality of VxVM.
<code>vol_kmsg_send_period</code>	This parameter cannot be tuned with the <code>vxtune</code> command. If specified in the <code>/kernel/drv/vxio.conf</code> file, the value is internally converted to microseconds. The value is then applied to the <code>vol_kmsg_send_period_usecs</code> parameter, which represents the value of the kernel message (KMSG) send period that is used by the clustering functionality of VxVM.

Table 39-2 Kernel tunable parameters for core VxVM (continued)

Parameter	Description
<code>vol_max_adminio_poolsz</code>	The maximum size of the memory pool that is used for administrative I/O operations. VxVM uses this pool when throttling administrative I/O. The default value is 64MB. The maximum size must not be greater than the value of the <code>voliorem_maxpool_sz</code> parameter.
<code>vol_max_vol</code>	This parameter cannot be tuned with the <code>vxtune</code> command. The maximum number of volumes that can be created on the system. The minimum permitted value is 1. The maximum permitted value is the maximum number of minor numbers representable on the system. The default value is 131071.
<code>vol_maxio</code>	The maximum size of logical I/O operations that can be performed without breaking up the request. I/O requests to VxVM that are larger than this value are broken up and performed synchronously. Physical I/O requests are broken up based on the capabilities of the disk device and are unaffected by changes to this maximum logical request limit. The default value is 2048 sectors (1MB). The value of <code>voliorem_maxpool_sz</code> must be at least 10 times greater than the value of <code>vol_maxio</code>. If DRL sequential logging is configured, the value of <code>voldrl_min_regionsz</code> must be set to at least half the value of <code>vol_maxio</code>.
<code>vol_maxioctl</code>	The maximum size of data that can be passed into VxVM via an <code>ioctl</code> call. Increasing this limit allows larger operations to be performed. Decreasing the limit is not generally recommended, because some utilities depend upon performing operations of a certain size and can fail unexpectedly if they issue oversized <code>ioctl</code> requests. The default value is 32768 bytes (32KB).

Table 39-2 Kernel tunable parameters for core VxVM (*continued*)

Parameter	Description
<code>vol_maxparallelio</code>	The number of I/O operations that the <code>vxconfigd</code> daemon is permitted to request from the kernel in a single <code>VOL_VOLDIO_READ</code> per <code>VOL_VOLDIO_WRITE</code> <code>ioctl</code> call. The default value is 256. This value should not be changed.
<code>vol_maxspecialio</code>	The maximum size of an I/O request that can be issued by an <code>ioctl</code> call. Although the <code>ioctl</code> request itself can be small, it can request that a large I/O request be performed. This tunable limits the size of these I/O requests. If necessary, a request that exceeds this value can be failed, or the request can be broken up and performed synchronously. The default value is 4096 sectors (2MB). Raising this limit can cause difficulties if the size of an I/O request causes the process to take more memory or kernel virtual mapping space than exists and thus deadlock. The maximum limit for this tunable is 20% of the smaller of physical memory or kernel virtual memory. It is inadvisable to go over this limit, because deadlock is likely to occur. If stripes are larger than the value of this tunable, full stripe I/O requests are broken up, which prevents full-stripe read/writes. This throttles the volume I/O throughput for sequential I/O or larger I/O requests. This tunable limits the size of an I/O request at a higher level in VxVM than the level of an individual disk. For example, for an 8 by 64KB stripe, a value of 256KB only allows I/O requests that use half the disks in the stripe; thus, it cuts potential throughput in half. If you have more columns or you have used a larger interleave factor, then your relative performance is worse. This tunable must be set, as a minimum, to the size of your largest stripe (RAID-0 or RAID-5).
<code>vol_stats_enable</code>	Enables or disables the I/O stat collection for Veritas Volume manager objects. The default value is 1, since this functionality is enabled by default.

Table 39-2 Kernel tunable parameters for core VxVM (continued)

Parameter	Description
<code>vol_subdisk_num</code>	The maximum number of subdisks that can be attached to a single plex. The default value of this tunable is 4096.
<code>voliomem_chunk_size</code>	The granularity of memory chunks used by VxVM when allocating or releasing system memory. A larger granularity reduces CPU overhead by allowing VxVM to retain hold of a larger amount of memory. The value of this tunable parameter depends on the page size of the system. You cannot specify a value larger than the default value. If you change the value, VxVM aligns the values to the page size when the system reboots. The default value is 524288 (512KB) for 8K page size.
<code>voliomem_maxpool_sz</code>	The maximum memory requested from the system by VxVM for internal purposes. This tunable has a direct impact on the performance of VxVM as it prevents one I/O operation from using all the memory in the system. VxVM allocates two pools that can grow up to this size, one for RAID-5 and one for mirrored volumes. Additional pools are allocated if instant (Copy On Write) snapshots are present. A write request to a RAID-5 volume that is greater than one fourth of the pool size is broken up and performed in chunks of one tenth of the pool size. A write request to a mirrored volume that is greater than the pool size is broken up and performed in chunks of the pool size. The minimum allowable value for this tunable is 5% of memory. If you specify a value less than that, the value is adjusted to 5% of value. The default value is 134217728 (128MB). The value of <code>voliomem_maxpool_sz</code> must be greater than the value of <code>volraid_minpool_size</code>. The value of <code>voliomem_maxpool_sz</code> must be at least 10 times greater than the value of <code>vol_maxio</code>.

Table 39-2 Kernel tunable parameters for core VxVM (continued)

Parameter	Description
<code>voliot_errbuf_dflt</code>	The default size of the buffer maintained for error tracing events. This buffer is allocated at driver load time and is not adjustable for size while VxVM is running. The default value is 16384 bytes (16KB). Increasing this buffer can provide storage for more error events at the expense of system memory. Decreasing the size of the buffer can result in an error not being detected via the tracing device. Applications that depend on error tracing to perform some responsive action are dependent on this buffer.
<code>voliot_iobuf_default</code>	The default size for the creation of a tracing buffer in the absence of any other specification of desired kernel buffer size as part of the trace <code>ioctl</code>. The default value is 8192 bytes (8KB). If trace data is often being lost due to this buffer size being too small, then this value can be increased.
<code>voliot_iobuf_limit</code>	The upper limit to the size of memory that can be used for storing tracing buffers in the kernel. Tracing buffers are used by the VxVM kernel to store the tracing event records. As trace buffers are requested to be stored in the kernel, the memory for them is drawn from this pool. Increasing this size can allow additional tracing to be performed at the expense of system memory usage. Setting this value to a size greater than can readily be accommodated on the system is inadvisable. The default value is 4194304 bytes (4MB).
<code>voliot_iobuf_max</code>	The maximum buffer size that can be used for a single trace buffer. Requests of a buffer larger than this size are silently truncated to this size. A request for a maximal buffer size from the tracing interface results (subject to limits of usage) in a buffer of this size. The default size for this buffer is 1048576 bytes (1MB). Increasing this buffer can provide for larger traces to be taken without loss for very heavily used volumes. Do not increase this value above the value for the <code>voliot_iobuf_limit</code> tunable value.

Table 39-2 Kernel tunable parameters for core VxVM (continued)

Parameter	Description
<code>voliot_max_open</code>	The maximum number of tracing channels that can be open simultaneously. Tracing channels are clone entry points into the tracing device driver. Each <code>vxtrace</code> process running on a system consumes a single trace channel. The default number of channels is 32. The allocation of each channel takes up approximately 20 bytes even when the channel is not in use.
<code>volraid_minpool_size</code>	This parameter cannot be tuned with the <code>vxtune</code> command. The initial amount of memory that is requested from the system by VxVM for RAID-5 operations. The maximum size of this memory pool is limited by the value of <code>voliomem_maxpool_sz</code>. The default value is 8192 sectors (4MB).
<code>volraid_rsrtransmax</code>	The maximum number of transient reconstruct operations that can be performed in parallel for RAID-5. A transient reconstruct operation is one that occurs on a non-degraded RAID-5 volume that has not been predicted. Limiting the number of these operations that can occur simultaneously removes the possibility of flooding the system with many reconstruct operations, and so reduces the risk of causing memory starvation. The default value is 1. Increasing this size improves the initial performance on the system when a failure first occurs and before a detach of a failing object is performed, but can lead to memory starvation.

Tunable parameters for FlashSnap (FMR)

Table 39-3 lists the kernel tunable parameters for FlashSnap. The `vxtune` command classifies these parameters under the FMR component.

You can tune the parameters using the `vxtune` command or the operating system method, unless otherwise noted.

Table 39-3 Kernel tunable parameters for FlashSnap (FMR)

Parameter	Description
vol_fmr_logsz	The maximum size in kilobytes of the bitmap that Non-Persistent FastResync uses to track changed blocks in a volume. The number of blocks in a volume that are mapped to each bit in the bitmap depends on the size of the volume, and this value changes if the size of the volume is changed. For example, if the volume size is 1 gigabyte and the system block size is 512 bytes, a value for this tunable of 4 yields a map that contains 16,384 bits, each bit representing one region of 128 blocks. The larger the bitmap size, the fewer the number of blocks that are mapped to each bit. This can reduce the amount of reading and writing required on resynchronization, at the expense of requiring more non-pageable kernel memory for the bitmap. Additionally, on clustered systems, a larger bitmap size increases the latency in I/O performance, and it also increases the load on the private network between the cluster members. This is because every other member of the cluster must be informed each time a bit in the map is marked. Since the region size must be the same on all nodes in a cluster for a shared volume, the value of this tunable on the master node overrides the tunable values on the slave nodes, if these values are different. Because the value of a shared volume can change, the value of this tunable is retained for the life of the volume. In configurations which have thousands of mirrors with attached snapshot plexes, the total memory overhead can represent a significantly higher overhead in memory consumption than is usual for VxVM. The default value is 4KB. The maximum and minimum permitted values are 1KB and 8KB. Specify a value to <code>vxtune</code> in Kilobytes. Note: The value of this tunable does not have any effect on Persistent FastResync.

Table 39-3 Kernel tunable parameters for FlashSnap (FMR) (continued)

Parameter	Description
<code>voldrl_dirty_regions</code>	This parameter applies to enhanced DCO layout (version 30) only. Represents the number of dirty regions to cache before another write to the same region causes a DRL update. A smaller number results in more frequent updates to the DRL, which decreases performance. A larger number results in better I/O performance, but requires that the DRL uses more memory. The default value is 1024.
<code>voldrl_max_drtregs</code>	The maximum number of dirty regions that can exist on the system for non-sequential DRL on volumes. A larger value may result in improved system performance at the expense of recovery time. This tunable can be used to regulate the worse-case recovery time for the system following a failure. The default value is 2048.
<code>voldrl_max_seq_dirty</code>	The maximum number of dirty regions allowed for sequential DRL. This is useful for volumes that are usually written to sequentially, such as database logs. Limiting the number of dirty regions allows for faster recovery if a crash occurs. The default value is 3.
<code>voldrl_min_regionsz</code>	The minimum number of sectors for a dirty region logging (DRL) volume region. With DRL, VxVM logically divides a volume into a set of consecutive regions. Larger region sizes tend to cause the cache hit-ratio for regions to improve. This improves the write performance, but it also prolongs the recovery time. The default value is 512 sectors. If DRL sequential logging is configured, the value of <code>voldrl_min_regionsz</code> must be set to at least half the value of <code>vol_maxio</code>. Specify the value in sectors.

Table 39-3 Kernel tunable parameters for FlashSnap (FMR) *(continued)*

Parameter	Description
<code>voldrl_volumemax_drtregs</code>	Maximum per-volume limit on dirty regions for a mirrored volume using traditional DRL. For heavily-used volumes, increase the value of this parameter to improve performance. The default value is 256.
<code>voldrl_volumemax_drtregs_20</code>	Maximum per-volume limit on dirty regions for a mirrored volume using version 20 DCO. For heavily-used volumes, increase the value of this parameter to improve performance. The default value is 1024.

Table 39-3 Kernel tunable parameters for FlashSnap (FMR) *(continued)*

Parameter	Description
volpagemod_max_memsz	

Table 39-3 Kernel tunable parameters for FlashSnap (FMR) (*continued*)

Parameter	Description
	The amount of memory that is allocated for caching FastResync and cache object metadata. The memory allocated for this cache is exclusively dedicated and is not available for other processes or applications. The default value is 65536KB (64MB). The valid range for this tunable is from 0 to 50% of physical memory. If cache objects or volumes that are prepared for instant snapshot operations are present on the system, setting the value below 512KB fails. If you do not use the FastResync or DRL features that are implemented using a version 20 DCO volume, you can set the value to 0. If you subsequently decide to enable these features, change the value to an appropriate one. Specify the value in kilobytes. The new value is page-aligned automatically; however the actual value specified is made persistent. Determine the value based on the region size and the number of volumes for which instant snapshots are taken. The paging module size must be at least twice the size required for the largest size volume, as calculated with the following formula: $size\_in\_KB = 6 * (total\_volume\_size\_in\_GB) * (64 / region\_size\_in\_KB)$ For example, a single 1TB volume requires around 6MB of paging memory if the region size is 64KB. The minimum value for the tunable parameter is at least twice that, or 12 MB. If there are multiple volumes, all volumes share the same paging module. The maximum requirement is calculated by multiplying the above formula by the number of volumes. However, a more reasonable value depends on the average load to each volume. For example, if only 20% of the data on each volume is updated, the paging module size can be reduced proportionally without compromising the performance. The minimum requirement for the largest volume still must be met. For example, if there are 10 volumes of 1TB each, the initial calculation is 60MB of paging

Table 39-3 Kernel tunable parameters for FlashSnap (FMR) *(continued)*

Parameter	Description
	memory. If only 20% of the data is updated, calculate the revised value as 12MB.

Tunable parameters for CVM

[Table 39-4](#) lists the kernel tunable parameter for CVM. This tunable cannot be tuned with the `vxtune` command.

You can tune the parameters using the `vxtune` command or the operating system method, unless otherwise noted.

Table 39-4 Kernel tunable parameter for CVM

<code>volcvm_smartsync</code>	If set to 0, <code>volcvm_smartsync</code> disables SmartSync on shared disk groups. If set to 1, this parameter enables the use of SmartSync with shared disk groups. The default value is 1. See “SmartSync recovery accelerator” on page 128.

Tunable parameters for VVR

[Table 39-5](#) lists the tunable parameters for VVR.

You can tune the parameters using the `vxtune` command or the operating system method, unless otherwise noted.

Table 39-5 VVR Tunables

Tunable Name	Description
<code>vol_cmpres_enabled</code>	A per-system tunable parameter that enables or disables compression globally. The default value is 0, since compression is disabled by default.
<code>vol_cmpres_threads</code>	A per-system tunable that lets you set the number of compression threads on the Primary host or the number of decompression threads on the Secondary host between 1 and 64. The default value is 10. You can tune this setting dependent on your CPU usage.

Table 39-5 VVR Tunables (continued)

Tunable Name	Description
<code>vol_dcm_replay_size</code>	This tunable cannot be changed using the <code>vxtune</code> command. The size of the DCM replay blocks. The default value is 256KB.
<code>vol_max_nmpool_sz</code>	The amount of buffer space available for requests coming in to the Secondary over the network. The default value is 64MB.
<code>vol_max_rdback_sz</code>	The amount of buffer space available for readbacks. The default value is 128MB.
<code>vol_max_wrspool_sz</code>	The write ship buffer space, which is the amount of buffer space that can be allocated on the logowner to receive writes sent by the non-logowner. The default value is 64MB.
<code>vol_min_lowmem_sz</code>	The minimum buffer space. VVR frees the write if the amount of buffer space available is below this threshold. The default value is about 520 kilobytes, twice the size of a maximum write. This value is auto-tunable. The value that you specify is used as an initial value and could change depending on the application write behavior.
<code>vol_nm_hb_timeout</code>	The heartbeat timeout value. The default value is 10 seconds.
<code>vol_rvio_maxpool_sz</code>	The amount of buffer space that can be allocated within the operating system to handle incoming writes. The default value is 128MB.
<code>vol_vvr_use_nat</code>	This tunable cannot be changed using the <code>vxtune</code> command. This tunable parameter directs VVR to use the translated address from the received message so that VVR can communicate over a NAT-based firewall. Set this tunable to 1 only if there is a NAT-based firewall in the configuration.

Points to note when changing the values of the VVR tunables

Note the following points when changing the values of the tunables:

- When decreasing the value of the `vol_rvio_maxpool_sz` tunable, all the RVGs on the host must be stopped.
- When decreasing the size of the tunables `vol_max_rdback_sz` and `vol_max_nmpool_sz` pause the RLINKs.
- The `vol_min_lowmem_sz` tunable is auto-tunable; depending on the incoming writes VVR increases or decreases the tunable value.
Auto-tuning is only supported for the tunable `vol_min_lowmem_sz`.

In a shared disk group environment, you may choose to set only those tunables that are required on each host. However, we recommend that you configure the tunables appropriately even if the tunables are currently not being used. This is because if the logowner changes, then tunables on the new logowner will be used. The following list of tunables are required to be set only on the logowner and not the other hosts:

- `vol_max_rdback_sz`
- `vol_max_nmpool_sz`
- `vol_max_wrspool_sz`
- `vol_dcm_replay_size`
- `vol_nm_hb_timeout`
- `vol_vvr_use_nat`

The tunable changes that are done using the `vxtune` command affect only the tunable values on the host on which it is run. Therefore, in a shared disk group environment, you must run the command separately on each host for which you want to change the tunable values.

How to change tunable parameters

Use the information in the following sections to tune the parameters for the Storage Foundation components.

Methods to change Veritas Dynamic Multi-Pathing tunable parameters

Veritas Dynamic Multi-Pathing (DMP) provides a variety of parameters that you can use to tune your configuration.

See [“DMP tunable parameters”](#) on page 884.

Change the DMP tunable parameters with one of the following methods:

Use the `vxddmpadm settune` command to display or modify the values.

See [“Changing the values of DMP parameters with the vxddmpadm settune command line”](#) on page 907.

Use the template method of the `vxddmpadm` command.

See [“About tuning Veritas Dynamic Multi-Pathing \(DMP\) with templates”](#) on page 907.

Changing the values of DMP parameters with the vxddmpadm settune command line

To set a DMP tunable parameter, use the following command:

```
# vxddmpadm settune dmp_tunable=value
```

To display the values of the DMP tunable parameters, use the following command:

```
# vxddmpadm gettune [dmp_tunable]
```

You can also use the template method to view or change DMP tunable parameters.

See [“About tuning Veritas Dynamic Multi-Pathing \(DMP\) with templates”](#) on page 907.

About tuning Veritas Dynamic Multi-Pathing (DMP) with templates

Veritas Dynamic Multi-Pathing has multiple tunable parameters and attributes that you can configure for optimal performance. In this release, DMP introduces a template method to update several tunable parameters and attributes with a single operation. The template represents a full or partial DMP configuration, showing the values of the parameters and attributes of the host.

To view and work with the tunable parameters, you can dump the configuration values of the DMP tunable parameters to a file. Edit the parameters and attributes, if required. Then, load the template file to a host to update all of the values in a single operation.

You can load the configuration file to the same host, or to another similar host. The template method is useful for the following scenarios:

- Configure multiple similar hosts with the optimal performance tuning values. Configure one host for optimal performance. After you have configured the host, dump the tunable parameters and attributes to a template file. You can then load the template file to another host with similar requirements. Symantec recommends that the hosts that use the same configuration template have the same operating system and similar I/O requirements.
- Define multiple specialized templates to handle different I/O load requirements.

When the load changes on a host, you can load a different template for the best performance. This strategy is appropriate for predictable, temporary changes in the I/O load. As the system administrator, after you define the system's I/O load behavior, you can customize tuning templates for particular loads. You can then automate the tuning, since there is a single load command that you can use in scripts or cron jobs.

At any time, you can reset the configuration, which reverts the values of the tunable parameters and attributes to the DMP default values.

You can manage the DMP configuration file with the `vxddmpadm config` commands.

See the `vxddmpadm(1m)` man page.

DMP tuning templates

The template mechanism enables you to tune DMP parameters and attributes by dumping the configuration values to a file, or to standard output.

DMP supports tuning the following types of information with template files:

- DMP tunable parameters.
- DMP attributes defined for an enclosure, array name, or array type.
- Veritas naming scheme parameters.

The template file is divided into sections, as follows:

DMP Tunables	Applied to all enclosures and arrays.
Namingscheme	Applied to all enclosures and arrays.
Arraytype	Use to customize array types. Applied to all of the enclosures of the specified array type.
Arrayname	Use if particular arrays need customization; that is, if the tunables vary from those applied for the array type. Attributes in this section are applied to all of the enclosures of the specified array name.
Enclosurename	Applied to the enclosures of the specified Cab serial number and array name. Use if particular enclosures need customization; that is, if the tunables vary from those applied for the array type and array name.

Loading is atomic for the section. DMP loads each section only if all of the attributes in the section are valid. When all sections have been processed, DMP reports the list of errors and warns the user. DMP does not support a partial rollback. DMP verifies the tunables and attributes during the load process. However, Symantec recommends that you check the configuration template file before you attempt to load the file. Make any required corrections until the configuration file validates correctly.

The attributes are given priority in the following order when a template is loaded:

Enclosure Section > Array Name Section > Array Type Section

If all enclosures of the same array type need the same settings, then remove the corresponding array name and enclosure name sections from the template. Define the settings only in the array type section. If some of the enclosures or array names need customized settings, retain the attribute sections for the array names or enclosures. You can remove the entries for the enclosures or the array names if they use the same settings that are defined for the array type.

When you dump a configuration file from a host, that host may contain some arrays which are not visible on the other hosts. When you load the template to a target host that does not include the enclosure, array type, or array name, DMP ignores the sections.

You may not want to apply settings to non-shared arrays or some host-specific arrays on the target hosts. Be sure to define an enclosure section for each of those arrays in the template. When you load the template file to the target host, the enclosure section determines the settings. Otherwise, DMP applies the settings from the respective array name or array type sections.

Example DMP tuning template

This section shows an example of a DMP tuning template.

DMP Tunables

```
dmp_cache_open=on
dmp_daemon_count=10
dmp_delayq_interval=15
dmp_restore_state=enabled
dmp_fast_recovery=on
dmp_health_time=60
dmp_log_level=1
dmp_low_impact_probe=on
dmp_lun_retry_timeout=0
dmp_path_age=300
dmp_pathswitch_blks_shift=9
dmp_probe_idle_lun=on
```

```

dmp_probe_threshold=5
dmp_restore_cycles=10
dmp_restore_interval=300
dmp_restore_policy=check_disabled
dmp_retry_count=5
dmp_scsi_timeout=30
dmp_sfg_threshold=1
dmp_stat_interval=1
dmp_monitor_ownership=on
dmp_native_multipathing=off
dmp_monitor_fabric=on
dmp_monitor_osevent=off
dmp_native_support=off

Namingscheme
    namingscheme=ebn
    persistence=yes
    lowercase=yes
    use_avid=yes

Arraytype
    arraytype=CLR-A/PF
    iopolicy=minimumq
    partitionsize=512
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0

Arraytype
    arraytype=ALUA
    iopolicy=adaptive
    partitionsize=512
    use_all_paths=no
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0

Arraytype
    arraytype=Disk
    iopolicy=minimumq
    partitionsize=512
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0

Arrayname
    arrayname=EMC_CLARiON
    iopolicy=minimumq

```

```

        partitionsize=512
        recoveryoption=nothrottle
        recoveryoption=timebound iotimeout=300
        redundancy=0
Arrayname
    arrayname=EVA4K6K
    iopolicy=adaptive
    partitionsize=512
    use_all_paths=no
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0
Arrayname
    arrayname=Disk
    iopolicy=minimumq
    partitionsize=512
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0
Enclosure
    serial=CK200051900278
    arrayname=EMC_CLARiion
    arraytype=CLR-A/PF
    iopolicy=minimumq
    partitionsize=512
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0
    dmp_lun_retry_timeout=0
Enclosure
    serial=50001FE1500A8F00
    arrayname=EVA4K6K
    arraytype=ALUA
    iopolicy=adaptive
    partitionsize=512
    use_all_paths=no
    recoveryoption=nothrottle
    recoveryoption=timebound iotimeout=300
    redundancy=0
    dmp_lun_retry_timeout=0
Enclosure
    serial=50001FE1500BB690
    arrayname=EVA4K6K

```

```

        arraytype=ALUA
        iopolicy=adaptive
        partitionsize=512
        use_all_paths=no
        recoveryoption=nothrottle
        recoveryoption=timebound iotimeout=300
        redundancy=0
        dmp_lun_retry_timeout=0
Enclosure
        serial=DISKS
        arrayname=Disk
        arraytype=Disk
        iopolicy=minimumq
        partitionsize=512
        recoveryoption=nothrottle
        recoveryoption=timebound iotimeout=300
        redundancy=0
        dmp_lun_retry_timeout=0

```

Tuning a DMP host with a configuration attribute template

You can use a template file to upload a series of changes to the DMP configuration to the same host or to another similar host.

In this release, Symantec recommends that you load the DMP template to a host that is similar to the host that was the source of the tunable values.

To configure DMP on a host with a template

- 1 Dump the contents of the current host configuration to a file.

```
# vxdmpadm config dump file=filename
```

- 2 Edit the file to make any required changes to the tunable parameters in the template.

The target host may include non-shared arrays or host-specific arrays. To avoid updating these with settings from the array name or array type, define an enclosure section for each of those arrays in the template. When you load the template file to the target host, the enclosure section determines the settings. Otherwise, DMP applies the settings from the respective array name or array type sections.

3 Validate the values of the DMP tunable parameters.

```
# vxddmpadm config check file=filename
```

DMP displays no output if the configuration check is successful. If the file contains errors, DMP displays the errors. Make any required corrections until the configuration file is valid. For example, you may see errors such as the following:

```
VxVM vxddmpadm ERROR V-5-1-0 Template file 'error.file' contains
following errors:
```

```
Line No: 22  'dmp_daemon_count' can not be set to 0 or less
Line No: 44  Specified value for 'dmp_health_time' contains
non-digits
Line No: 64  Specified value for 'dmp_path_age' is beyond
the limit of its value
Line No: 76  'dmp_probe_idle_lun' can be set to either on or off
Line No: 281 Unknown arraytype
```

4 Load the file to the target host.

```
# vxddmpadm config load file=filename
```

During the loading process, DMP validates each section of the template. DMP loads all valid sections. DMP does not load any section that contains errors.

Managing the DMP configuration files

You can display the name of the template file most recently loaded to the host. The information includes the date and time when DMP loaded the template file.

To display the name of the template file that the host currently uses

```
◆ # vxddmpadm config show
```

```
TEMPLATE_FILE      DATE              TIME
=====
/tmp/myconfig      Feb 09, 2011     11:28:59
```

Resetting the DMP tunable parameters and attributes to the default values

DMP maintains the default values for the DMP tunable parameters and attributes. At any time, you can restore the default values to the host. Any changes that you applied to the host with template files are discarded.

To reset the DMP tunables to the default values

- ◆ Use the following command:

```
# vxddmpadm config reset
```

DMP tunable parameters and attributes that are supported for templates

DMP supports tuning the following tunable parameters and attributes with a configuration template.

DMP tunable parameters

See [“DMP tunable parameters”](#) on page 884.

DMP attributes defined for an enclosure, array name, or array type.

- iopolicy
- partitionsize
- use_all_paths
- recoveryoption attributes (retrycount or iotimeout)
- redundancy
- dmp_lun_retry_timeout

Naming scheme attributes:

- naming scheme
- persistence
- lowercase
- use_avid

The following tunable parameters are NOT supported with templates:

- OS tunables
- TPD mode
- Failover attributes of enclosures (failovermode)

Methods to change Veritas Volume Manager tunable parameters

Veritas Volume Manager (VxVM) provides a variety of parameters that you can use to tune your configuration.

See [“Tunable parameters for VxVM”](#) on page 892.

Change the VxVM tunable parameters with one of the following methods:

Use the `vxvtune` command to display or modify the values of the VxVM tunable parameters.

See [“Changing the values of the Veritas Volume Manager tunable parameters using the vxvtune command line”](#) on page 917.

Use the template method of the `vxtune` command.

See [“Changing the value of the Veritas Volume Manager tunable parameters using templates”](#) on page 920.

Use the `vxio.conf` file.

See [“Changing the tunable values using the `vxio.conf` file”](#) on page 921.

Displaying the values of the VxVM tunables

The `vxtune` command allows you to display the values for tunable parameters in the following ways:

- All of the VxVM tunable parameters
- Tunable parameters by component
- Single, specified parameter

To display all of the VxVM tunable parameters

- ◆ To display all of the VxVM the tunable parameters, use the `vxtune` command with no argument.

```
# vxtune [-u {h|unit}] [-lr]
```

The `-l` option displays a description along with the values in default units.

The `-r` option displays the tunable values in raw bytes.

The `-uh` option displays output in human-friendly units.

The Reboot column displays `N` if the value can be changed online. The column displays `Y` if a system reboot is required before the new value takes effect.

The following example shows a truncated output, that shows the format.

```
# vxtune
```

Tunable	Current Value	Default Value	Reboot
-----	-----	-----	-----
vol_checkpoint_default	20480	20480	Y
vol_cmpres_enabled	0	0	N
vol_cmpres_threads	10	10	N
vol_default_iodelay	50	50	Y
vol_fmr_logsz	4	4	Y
vol_max_adminio_poolsz	67108864	67108864	Y
vol_max_nmpool_sz	67108864	67108864	N
.			
.			
.			

To display the tunables for a VxVM component

- 1 The VxVM tunable parameters are classified into components.

To check the list of recognized components, use the following command:

```
# vxtune help components
```

- 2 To display the tunables for a VxVM component, specify the component name.

```
# vxtune component
```

For example, to display the supported tunables for VVR, use the following command:

```
# vxtune vvr
```

Tunable	Current Value	Default Value	Reboot
vol_cmpres_enabled	0	0	N
vol_cmpres_threads	10	10	N
vol_max_nmpool_sz	67108864	67108864	N
vol_max_rdback_sz	134217728	134217728	N
vol_max_wrspool_sz	67108864	67108864	N
vol_min_lowmem_sz	540672	540672	N
vol_nm_hb_timeout	10	10	N
vol_rvio_maxpool_sz	105224192	134217728	N

Additional options are -l to display a description, -r to display the output in bytes, and -uh to display the output in human-friendly units.

To display the value of a specific tunable

- ◆ To display the value of the specific tunable, use the following `vxtune` command on the required host.

```
# vxtune tunable_name
```

For example, to view the value of the tunable `vol_rvio_maxpool_sz`, use the following command:

```
# vxtune vol_rvio_maxpool_sz
Tunable                Current Value Default Value Reboot
-----
vol_rvio_maxpool_sz    105224192    134217728    N
```

Additional options are `-l` to display a description, `-r` to display the output in bytes, and `-h` to display the output in human-friendly units.

To display the output in bytes, use the `vxtune` command with the `-r` option.

```
# vxtune -r tunable_name
```

For example, to view the value of the tunable `vol_max_wrspool_sz` in bytes, use the following command:

```
# vxtune -r vol_max_wrspool_sz
Tunable                Current Value Default Value Reboot
-----
vol_max_wrspool_sz    67108864    67108864    N
```

Changing the values of the Veritas Volume Manager tunable parameters using the `vxtune` command line

Use the `vxtune` command to display or change the values of the VxVM tunable parameters. The changes are persistent so that the value is retained after subsequent reboots. Before setting a new value, VxVM validates the value to ensure that it is within the permissible range for the tunable. If the value is valid, VxVM updates the tunable. Some tunables require a reboot before the changed value takes effect. VxVM prompts you to reboot the system, if required.

VxVM updates the tunable values in the `vxio.conf` file.

For most tunables, specify the value of the tunable with a suffix to indicate the units: K, M, or G. If no unit is specified, `vxtune` assumes the value is bytes.

Note: The default unit for entering a value may differ from the default display unit.

See [“Displaying the values of the VxVM tunables”](#) on page 915.

To change the values of the VxVM tunable parameters

- 1 Find the name and current value of the tunable you want to change. Use the `-l` option to display a description.

```
# vxtune -l
```

The following example shows a truncated output, that shows the format.

Tunable	Current Value	Default Value	Reboot	Description
vol_checkpoint_default	20480	20480	Y	Size of VxVM checkpoints (sectors)
vol_cmpres_enabled	0	0	N	Allow enabling compression for VERITAS V
vol_cmpres_threads	10	10	N	Maximum number of compression threads for
vol_default_iodelay	50	50	Y	Time to pause between I/O requests from V
vol_fmr_logsz	4	4	Y	Maximum size of bitmap Fast Mirror Resync
vol_max_adminio_poolsz	67108864	67108864	Y	Maximum amount of memory used by VxVM adm
.				
.				
.				

The output displays the default value and the current value. The Reboot field indicates whether or not a reboot is required before the tunable value takes effect.

See the `vxtune(1M)` manual page.

- 2 Set the new value for a specific tunable. Specify the value with a suffix to indicate the units: K, M, or G. If not unit is specified, the `vxtune` command uses the default unit for the tunable parameter. For most tunables, the default value is bytes. The description in the `vxtune` output displays the default units for each tunable.

```
# vxtune tunable_name tunable_value
```

For example, to change the value of `vol_cmpres_enabled` to 1, use the following command:

```
# vxtune vol_cmpres_enabled 1
```

For example, to change the default value of the `vol_rvio_maxpool_sz` tunable to 160MB, use the following command.

```
# vxtune vol_rvio_maxpool_sz 160M
```

3 Verify the new value.

```
# vxtune tunable_name
```

For example, to view the changed value for `vol_cmpres_enabled`, use the following command:

```
# vxtune vol_cmpres_enabled
Tunable                Current Value Default Value Reboot
-----
vol_cmpres_enabled     1                0                N
```

For example, to view the changed value for the `vol_rvio_maxpool_sz` tunable parameter, use the following command:

```
# vxtune vol_rvio_maxpool_sz
Tunable                Current Value Default Value Reboot
-----
vol_rvio_maxpool_sz    167772160      134217728      N
```

Changing the value of the Veritas Volume Manager tunable parameters using templates

VxVM provides a template method to change the tunable parameters. With this method, you export the tunable parameters to a file, modify the file, then import the file. The tunable template must be strictly of the format that export provides. In case of discrepancies observe that that particular value will be discarded.

To change the values of the VxVM tunable parameters using templates

- 1 Export the tunable parameters and their values to a tunable template file. You can export all of the tunable parameters or specify a component.

```
# vxtune export file=file_name [component]
```

For example:

```
# vxtune export file=vxvm-tunables
```

```
# vxtune export file=vvr-tunables vvr
```

- 2 Modify the template as required. You must retain the file format that the export operation provides.
- 3 Import the tunable template file to the system. The import operation only applies valid values. If a value is not valid for a specific parameter, that particular value is discarded.

```
# vxtune import file=file_name
```

For example:

```
# vxtune import file=vxvm-tunables
```

Changing the tunable values using the vxio.conf file

You can tune the VxVM parameters by editing the `/kernel/drv/vxio.conf` file. The changes made to the tunables using this method are persistent, after a system reboot.

Warning: Do not edit the configuration file for the `vxspec` driver, `/kernel/drv/vxspec.conf`.

To change the value of a tunable

- 1 Display the current values of the tunables.

```
# prtconf -vP
.
.
.
vxio, instance #0
  System properties:
    name <vol_tunable> length <4>
    value <0x00001388>
  Driver properties:
    name <voldrl_max_seq_dirty> length <4>
    value <0x00000003>
.
.
.
```

All VxVM tunables that you specify in `/kernel/drv/vxio.conf` are listed in the output under the “System properties.” heading for the `vxio` drivers. All unchanged tunables are listed with their default values under the “Driver properties” heading. The sample output shows the new value for `vol_tunable` in hexadecimal:

For more information, see the `prtconf(1M)` and `driver.conf(4)` manual pages.

- 2 Navigate to the `/kernel/drv/` directory which contains the `vxio.conf` file.
- 3 Make a backup copy of the `/kernel/drv/vxio.conf` file.
- 4 Open the file using any editor. Add or edit the tunable in the `/kernel/drv/vxio.conf` file, using the following format:

```
tunable_name=value;
```

Where the *value* is specified in bytes.

For example, to change the value of the tunable `vol_rvio_maxpool_sz` to 128M, append the following line to the file `/kernel/drv/vxio.conf`:

```
vol_rvio_maxpool_sz=134217728;
```

The specified value for `vol_rvio_maxpool_sz` is now applicable system-wide.

- 5 Reboot the system for the changes to take effect.

Quotas

This chapter includes the following topics:

- [About quota limits](#)
- [About quota files on Veritas File System](#)
- [About quota commands](#)
- [About quota checking with Veritas File System](#)
- [Using quotas](#)

About quota limits

Veritas File System (VxFS) supports user and group quotas. The quota system limits the use of two principal resources of a file system: files and data blocks. For each of these resources, you can assign quotas to individual users and groups to limit their usage.

You can set the following kinds of limits for each of the two resources:

hard limit	An absolute limit that cannot be exceeded under any circumstances.
soft limit	Must be lower than the hard limit, and can be exceeded, but only for a limited time. The time limit can be configured on a per-file system basis only. The VxFS default limit is seven days.

Soft limits are typically used when a user must run an application that could generate large temporary files. In this case, you can allow the user to exceed the quota limit for a limited time. No allocations are allowed after the expiration of the time limit. Use the `vxedquota` command to set limits.

See [“Using quotas”](#) on page 926.

Although file and data block limits can be set individually for each user and group, the time limits apply to the file system as a whole. The quota limit information is associated with user and group IDs and is stored in a user or group quota file.

See [“About quota files on Veritas File System”](#) on page 924.

The quota soft limit can be exceeded when VxFS preallocates space to a file.

See [“About extent attributes”](#) on page 235.

About quota files on Veritas File System

A `quotas` file (named `quotas`) must exist in the root directory of a file system for any of the quota commands to work. For group quotas to work, there must be a `quotas.grp` file. The files in the file system's mount point are referred to as the external quotas file. VxFS also maintains an internal quotas file for its own use.

The quota administration commands read and write to the external quotas file to obtain or change usage limits. VxFS uses the internal file to maintain counts of data blocks and inodes used by each user. When quotas are turned on, the quota limits are copied from the external quotas file into the internal quotas file. While quotas are on, all the changes in the usage information and changes to quotas are registered in the internal quotas file. When quotas are turned off, the contents of the internal quotas file are copied into the external quotas file so that all data between the two files is synchronized.

VxFS supports group quotas in addition to user quotas. Just as user quotas limit file system resource (disk blocks and the number of inodes) usage on individual users, group quotas specify and limit resource usage on a group basis. As with user quotas, group quotas provide a soft and hard limit for file system resources. If both user and group quotas are enabled, resource utilization is based on the most restrictive of the two limits for a given user.

To distinguish between group and user quotas, VxFS quota commands use a `-g` and `-u` option. The default is user quotas if neither option is specified. One exception to this rule is when you specify the `-o quota` option as a `mount` command option. In this case, both user and group quotas are enabled. Support for group quotas also requires a separate group quotas file. The VxFS group quota file is named `quotas.grp`. The VxFS user quotas file is named `quotas`. This name was used to distinguish it from the `quotas.user` file used by other file systems under Solaris.

About quota commands

In general, quota administration for VxFS is performed using commands similar to UFS quota commands. On Solaris, the available quota commands are UFS-specific.

That is, these commands work only on UFS file systems. For this reason, VxFS supports a similar set of commands that work only for VxFS file systems.

Note: Most of the quota commands in VxFS are similar to BSD quota commands. However, the `quotacheck` command is an exception; VxFS does not support an equivalent command.

See [“About quota checking with Veritas File System”](#) on page 926.

VxFS supports the following quota-related commands:

<code>vxedquota</code>	Edits quota limits for users and groups. The limit changes made by <code>vxedquota</code> are reflected both in the internal quotas file and the external quotas file.
<code>vxrepquota</code>	Provides a summary of quotas and disk usage.
<code>vxquot</code>	Provides file ownership and usage summaries.
<code>vxquota</code>	Views quota limits and usage.
<code>vxquotaon</code>	Turns quotas on for a mounted VxFS file system.
<code>vxquotaoff</code>	Turns quotas off for a mounted VxFS file system.

The `vxquota`, `vxrepquota`, `vxquot`, and `vxedquota` commands support the `-H` option for human-friendly input and output. When the `-H` option is used, the storage size is displayed in the following human-friendly units: bytes (B), kilobytes (KB), megabytes (MB), gigabytes (GB), terabyte (TB), petabytes (PB), and exabytes (EB). The quota soft and hard limits, quota usage, and the total storage consumed by a specific user or group or all users or groups can be obtained in human-friendly units using the `-H` option.

Beside these commands, the VxFS `mount` command supports special mounts option (`-o quota|usrquota|grpquota`), which can be used to turn on quotas at mount time.

For additional information on the quota commands, see the corresponding manual pages.

Note: When VxFS file systems are exported via NFS, the VxFS quota commands on the NFS client cannot query or edit quotas. You can use the VxFS quota commands on the server to query or edit quotas.

About quota checking with Veritas File System

The standard practice with most quota implementations is to mount all file systems and then run a quota check on each one. The quota check reads all the inodes on disk and calculates the usage for each user and group. This can be time consuming, and because the file system is mounted, the usage can change while quotacheck is running.

VxFS does not support a `quotacheck` command. With VxFS, quota checking is performed automatically (if necessary) at the time quotas are turned on. A quota check is necessary if the file system has changed with respect to the usage information as recorded in the internal quotas file. This happens only if the file system was written with quotas turned off, or if there was structural damage to the file system that required a full file system check.

See the `fsck_vxfs(1M)` manual page.

A quota check generally reads information for each inode on disk and rebuilds the internal quotas file. It is possible that while quotas were not on, quota limits were changed by the system administrator. These changes are stored in the external quotas file. As part of enabling quotas processing, quota limits are read from the external quotas file into the internal quotas file.

Using quotas

The VxFS quota commands are used to perform the following quota functions:

- [Turning on quotas](#)
- [Turning on quotas at mount time](#)
- [Editing user and group quotas](#)
- [Modifying time limits](#)
- [Viewing disk quotas and usage](#)
- [Displaying blocks owned by users or groups](#)
- [Turning off quotas](#)

Turning on quotas

To use the quota functionality on a file system, quotas must be turned on. You can turn quotas on at mount time or after a file system is mounted.

Note: Before turning on quotas, the root directory of the file system must contain a file for user quotas named `quotas`, and a file for group quotas named `quotas.grp` owned by root.

To turn on quotas

- 1 To turn on user and group quotas for a VxFS file system, enter:

```
# vxquotaon /mount_point
```

- 2 To turn on only user quotas for a VxFS file system, enter:

```
# vxquotaon -u /mount_point
```

- 3 To turn on only group quotas for a VxFS file system, enter:

```
# vxquotaon -g /mount_point
```

Turning on quotas at mount time

Quotas can be turned on with the `mount` command when you mount a file system.

To turn on quotas at mount time

- 1 To turn on user or group quotas for a file system at mount time, enter:

```
# mount -F vxfs -o quota special /mount_point
```

- 2 To turn on only user quotas, enter:

```
# mount -F vxfs -o usrquota special /mount_point
```

- 3 To turn on only group quotas, enter:

```
# mount -F vxfs -o grpquota special  
/mount_point
```

Editing user and group quotas

You can set up user and group quotas using the `vxedquota` command. You must have superuser privileges to edit quotas.

`vxedquota` creates a temporary file for the given user; this file contains on-disk quotas for each mounted file system that has a `quotas` file. It is not necessary that

quotas be turned on for `vxedquota` to work. However, the quota limits are applicable only after quotas are turned on for a given file system.

To edit quotas

- 1 Specify the `-u` option to edit the quotas of one or more users specified by *username*:

```
# vxedquota [-u] username
```

Editing the quotas of one or more users is the default behavior if the `-u` option is not specified.

- 2 Specify the `-g` option to edit the quotas of one or more groups specified by *groupname*:

```
# vxedquota -g groupname
```

Modifying time limits

The soft and hard limits can be modified or assigned values. For any user or group, usage can never exceed the hard limit after quotas are turned on.

Modified time limits apply to the entire file system and cannot be set selectively for each user or group.

To modify time limits

- 1 Specify the `-t` option to modify time limits for any user:

```
# vxedquota [-u] -t
```

- 2 Specify the `-g` and `-t` options to modify time limits for any group:

```
# vxedquota -g -t
```

Viewing disk quotas and usage

Use the `vxquota` command to view a user's or group's disk quotas and usage on VxFS file systems.

To display disk quotas and usage

- 1 To display a user's quotas and disk usage on all mounted VxFS file systems where the `quotas` file exists, enter:

```
# vxquota -v [-u] username
```

- 2 To display a group's quotas and disk usage on all mounted VxFS file systems where the `quotas.grp` file exists, enter:

```
# vxquota -v -g groupname
```

Displaying blocks owned by users or groups

Use the `vxquot` command to display the number of blocks owned by each user or group in a file system.

To display the number of blocks owned by users or groups

- 1 To display the number of files and the space owned by each user, enter:

```
# vxquot [-u] -f filesystem
```

- 2 To display the number of files and the space owned by each group, enter:

```
# vxquot -g -f filesystem
```

Turning off quotas

Use the `vxquotaoff` command to turn off quotas.

To turn off quotas

- 1 To turn off quotas for a mounted file system, enter:

```
# vxquotaoff /mount_point
```

- 2 To turn off only user quotas for a VxFS file system, enter:

```
# vxquotaoff -u /mount_point
```

- 3 To turn off only group quotas for a VxFS file system, enter:

```
# vxquotaoff -g /mount_point
```

File Change Log

This chapter includes the following topics:

- [About File Change Log](#)
- [About the File Change Log file](#)
- [File Change Log administrative interface](#)
- [File Change Log programmatic interface](#)
- [Summary of API functions](#)

About File Change Log

The VxFS File Change Log (FCL) tracks changes to files and directories in a file system.

Applications that typically use the FCL are usually required to:

- scan an entire file system or a subset
- discover changes since the last scan

These applications may include: backup utilities, webcrawlers, search engines, and replication programs.

Note: The FCL tracks when the data has changed and records the change type, but does not track the actual data changes. It is the responsibility of the application to examine the files to determine the changed data.

FCL functionality is a separately licensable feature.

See the *Veritas Storage Foundation Cluster File System High Availability Release Notes*.

About the File Change Log file

File Change Log records file system changes such as creates, links, unlinks, renaming, data appended, data overwritten, data truncated, extended attribute modifications, holes punched, and miscellaneous file property updates.

FCL stores changes in a sparse file in the file system namespace. The FCL file is located in `mount_point/lost+found/changelog`. The FCL file behaves like a regular file, but some operations are prohibited. The standard system calls `open(2)`, `lseek(2)`, `read(2)` and `close(2)` can access the data in the FCL, while the `write(2)`, `mmap(2)` and `rename(2)` calls are not allowed.

Warning: Although some standard system calls are currently supported, the FCL file might be pulled out of the namespace in future VxFS release and these system calls may no longer work. It is recommended that all new applications be developed using the programmatic interface.

The FCL log file contains both the information about the FCL, which is stored in the FCL superblock, and the changes to files and directories in the file system, which is stored as FCL records.

See [“File Change Log programmatic interface”](#) on page 934.

In the 4.1 release, the structure of the File Change Log file was exposed through the `/opt/VRTS/include/sys/fs/fcl.h` header file. In this release, the internal structure of the FCL file is opaque. The recommended mechanism to access the FCL is through the API described by the `/opt/VRTSfssdk/5.1.100.000/include/vxfsutil.h` header file.

The `/opt/VRTS/include/sys/fs/fcl.h` header file is included in this release to ensure that applications accessing the FCL with the 4.1 header file do not break. New applications should use the new FCL API described in `/opt/VRTSfssdk/6.0.000.000/include/vxfsutil.h`. Existing applications should also be modified to use the new FCL API.

To provide backward compatibility for the existing applications, this release supports multiple FCL versions. Users have the flexibility of specifying the FCL version for new FCLs. The default FCL version is 4.

See the `fcladm(1M)` man page.

File Change Log administrative interface

The FCL can be set up and tuned through the `fcladm` and `vxtunefs VxFS` administrative commands.

See the `fcladm(1M)` and `vxtunefs(1M)` manual pages.

The FCL keywords for `fcladm` are as follows:

<code>clear</code>	Disables the recording of the audit, open, close, and statistical events after it has been set.
<code>dump</code>	Creates a regular file image of the FCL file that can be downloaded to an off-host processing system. This file has a different format than the FCL file.
<code>on</code>	Activates the FCL on a mounted file system. VxFS 5.0 and later releases support either FCL Versions 3 or 4. If no version is specified, the default is Version 4. Use <code>fcladm on</code> to specify the version.
<code>print</code>	Prints the contents of the FCL file starting from the specified offset.
<code>restore</code>	Restores the FCL file from the regular file image of the FCL file created by the <code>dump</code> keyword.
<code>rm</code>	Removes the FCL file. You must first deactivate the FCL with the <code>off</code> keyword, before you can remove the FCL file.
<code>set</code>	Enables the recording of events specified by the 'eventlist' option. See the <code>fcladm(1M)</code> manual page.
<code>state</code>	Writes the current state of the FCL to the standard output.
<code>sync</code>	Brings the FCL to a stable state by flushing the associated data of an FCL recording interval.

The FCL tunable parameters for `vxtunefs` are as follows:

<code>fcl_keeptime</code>	Specifies the duration in seconds that FCL records stay in the FCL file before they can be purged. The first records to be purged are the oldest ones, which are located at the beginning of the file. Additionally, records at the beginning of the file can be purged if allocation to the FCL file exceeds <code>fcl_maxalloc</code> bytes. The default value of <code>fcl_keeptime</code> is 0. If the <code>fcl_maxalloc</code> parameter is set, records are purged from the FCL file if the amount of space allocated to the FCL file exceeds <code>fcl_maxalloc</code> . This is true even if the elapsed time the records have been in the log is less than the value of <code>fcl_keeptime</code> .

<code>fcl_maxalloc</code>	Specifies the maximum number of spaces in bytes to be allocated to the FCL file. When the space allocated exceeds <code>fcl_maxalloc</code>, a hole is punched at the beginning of the file. As a result, records are purged and the first valid offset (<code>fc_foff</code>) is updated. In addition, <code>fcl_maxalloc</code> may be violated if the oldest record has not reached <code>fcl_keeptime</code>. The minimum value of <code>fcl_maxalloc</code> is 4 MB. The default value is <code>fs_size/33</code>.
<code>fcl_winterval</code>	Specifies the time in seconds that must elapse before the FCL records an overwrite, extending write, or a truncate. This helps to reduce the number of repetitive records in the FCL. The <code>fcl_winterval</code> timeout is per inode. If an inode happens to go out of cache and returns, its write interval is reset. As a result, there could be more than one write record for that file in the same write interval. The default value is 3600 seconds.
<code>fcl_ointerval</code>	The time interval in seconds within which subsequent opens of a file do not produce an additional FCL record. This helps to reduce the number of repetitive records logged in the FCL file. If the tracking of access information is also enabled, a subsequent file open even within the <code>fcl_ointerval</code> may produce a record, if it is opened by a different user. Similarly, if the inode is bumped out of cache, this may also produce more than one record within the same open interval. The default value is 600 sec.

Either or both `fcl_maxalloc` and `fcl_keeptime` must be set to activate the FCL feature. The following are examples of using the `fcladm` command.

To activate FCL for a mounted file system, type the following:

```
# fcladm on mount_point
```

To deactivate the FCL for a mounted file system, type the following:

```
# fcladm off mount_point
```

To remove the FCL file for a mounted file system, on which FCL must be turned off, type the following:

```
# fcladm rm mount_point
```

To obtain the current FCL state for a mounted file system, type the following:

```
# fcladm state mount_point
```

To enable tracking of the file opens along with access information with each event in the FCL, type the following:

```
# fcladm set fileopen,accessinfo mount_point
```

To stop tracking file I/O statistics in the FCL, type the following:

```
# fcladm clear filestats mount_point
```

Print the on-disk FCL super-block in text format to obtain information about the FCL file by using offset 0. Because the FCL on-disk super-block occupies the first block of the FCL file, the first and last valid offsets into the FCL file can be determined by reading the FCL super-block and checking the `fc_off` field. Enter:

```
# fcladm print 0 mount_point
```

To print the contents of the FCL in text format, of which the offset used must be 32-byte aligned, enter:

```
# fcladm print offset mount_point
```

File Change Log programmatic interface

VxFS provides an enhanced API to simplify reading and parsing the FCL file in two ways:

Simplified reading	The API simplifies user tasks by reducing additional code needed to parse FCL file entries. In 4.1, to obtain event information such as a remove or link, the user was required to write additional code to get the name of the removed or linked file. In this release, the API allows the user to directly read an assembled record. The API also allows the user to specify a filter to indicate a subset of the event records of interest.
Backward compatibility	Providing API access for the FCL feature allows backward compatibility for applications. The API allows applications to parse the FCL file independent of the FCL layout changes. Even if the hidden disk layout of the FCL changes, the API automatically translates the returned data to match the expected output record. As a result, the user does not need to modify or recompile the application due to changes in the on-disk FCL layout.

The following sample code fragment reads the FCL superblock, checks that the state of the FCL is `VX_FCLS_ON`, issues a call to `vxfs_fcl_sync` to obtain a finishing offset to read to, determines the first valid offset in the FCL file, then reads the

entries in 8K chunks from this offset. The section process fcl entries is what an application developer must supply to process the entries in the FCL file.

```
#include <stdint.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/fcntl.h>
#include <errno.h>
#include <fcl.h>
#include <vxfsutil.h>
#define FCL_READSZ 8192
char* fclname = "/mnt/lost+found/changelog";
int read_fcl(fclname) char* fclname;
{
    struct fcl_sb fclsb;
    uint64_t off, lastoff;
    size_t size;
    char buf[FCL_READSZ], *bufp = buf;
    int fd;
    int err = 0;
    if ((fd = open(fclname, O_RDONLY)) < 0) {
        return ENOENT;
    }
    if ((off = lseek(fd, 0, SEEK_SET)) != 0) {
        close(fd);
        return EIO;
    }
    size = read(fd, &fclsb, sizeof (struct fcl_sb));
    if (size < 0) {
        close(fd);
        return EIO;
    }
    if (fclsb.fc_state == VX_FCLS_OFF) {
        close(fd);
        return 0;
    }
    if (err = vxfs_fcl_sync(fclname, &lastoff)) {
        close(fd);
        return err;
    }
    if ((off = lseek(fd, off_t, uint64_t)) != uint64_t) {
        close(fd);
    }
}
```

```

        return EIO;
    }
    while (off < lastoff) {
        if ((size = read(fd, bufp, FCL_READSZ)) <= 0) {
            close(fd);
            return errno;
        }
        /* process fcl entries */
        off += size;
    }
    close(fd);
    return 0;
}

```

Summary of API functions

The following is a brief summary of File Change Log API functions:

<code>vxfs_fcl_close()</code>	Closes the FCL file and cleans up resources associated with the handle.
<code>vxfs_fcl_cookie()</code>	Returns an opaque structure that embeds the current FCL activation time and the current offset. This cookie can be saved and later passed to <code>vxfs_fcl_seek()</code> function to continue reading from where the application last stopped.
<code>vxfs_fcl_getinfo()</code>	Returns information such as the state and version of the FCL file.
<code>vxfs_fcl_open()</code>	Opens the FCL file and returns a handle that can be used for further operations.
<code>vxfs_fcl_read()</code>	Reads FCL records of interest into a buffer specified by the user.
<code>vxfs_fcl_seek()</code>	Extracts data from the specified cookie and then seeks to the specified offset.
<code>vxfs_fcl_seektime()</code>	Seeks to the first record in the FCL after the specified time.

Reverse path name lookup

This chapter includes the following topics:

- [Reverse path name lookup](#)

Reverse path name lookup

The reverse path name lookup feature obtains the full path name of a file or directory from the inode number of that file or directory. The inode number is provided as an argument to the `vxlsino` administrative command, or the `vxfs_inotopath_gen(3)` application programming interface library function.

The reverse path name lookup feature can be useful for a variety of applications, such as for clients of the VxFS File Change Log feature, in backup and restore utilities, and for replication products. Typically, these applications store information by inode numbers because a path name for a file or directory can be very long, thus the need for an easy method of obtaining a path name.

An inode is a unique identification number for each file in a file system. An inode contains the data and metadata associated with that file, but does not include the file name to which the inode corresponds. It is therefore relatively difficult to determine the name of a file from an inode number. The `ncheck` command provides a mechanism for obtaining a file name from an inode identifier by scanning each directory in the file system, but this process can take a long period of time. The VxFS reverse path name lookup feature obtains path names relatively quickly.

Note: Because symbolic links do not constitute a path to the file, the reverse path name lookup feature cannot track symbolic links to files.

Because of the possibility of errors with processes renaming or unlinking and creating new files, it is advisable to perform a `lookup` or `open` with the path name and verify that the inode number matches the path names obtained.

See the `vxlsino(1M)`, `vxfs_inotopath_gen(3)`, and `vxfs_inotopath(3)` manual pages.

Disk layout

This appendix includes the following topics:

- [About disk layouts](#)
- [VxFS Version 7 disk layout](#)
- [VxFS Version 8 disk layout](#)
- [VxFS Version 9 disk layout](#)
- [Using UNIX Commands on File Systems Larger than One TB](#)

About disk layouts

The disk layout is the way file system information is stored on disk. On VxFS, several different disk layout versions were created to take advantage of evolving technological developments.

The disk layout versions used on VxFS are:

Version 1	Version 1 disk layout is the original VxFS disk layout provided with pre-2.0 versions of VxFS.	Not Supported
Version 2	Version 2 disk layout supports features such as filesets, dynamic inode allocation, and enhanced security. The Version 2 layout is available with and without quotas support.	Not Supported
Version 3	Version 3 disk layout encompasses all file system structural information in files, rather than at fixed locations on disk, allowing for greater scalability. Version 3 supports files and file systems up to one terabyte in size.	Not Supported

Version 4	Version 4 disk layout encompasses all file system structural information in files, rather than at fixed locations on disk, allowing for greater scalability. Version 4 supports files and file systems up to one terabyte in size.	Not Supported
Version 5	Version 5 enables the creation of file system sizes up to 32 terabytes. File sizes can be a maximum of 4 billion file system blocks. File systems larger than 1TB must be created on a Veritas Volume Manager volume.	Not Supported
Version 6	Version 6 disk layout enables features such as multi-volume support, cross-platform data sharing, named data streams, and File Change Log. A disk layout Version 6 file system can still be mounted, but this will be disallowed in future releases. Symantec recommends that you upgrade from Version 6 to the latest default disk layout version. In this release, disk layout Version 6 cannot be cluster mounted. You cannot create new file systems with disk layout Version 6. The only operation that you can perform on a file system with disk layout Version 6 is to upgrade the disk layout to a supported version. If you upgrade a file system from disk layout Version 6 to a later version, once the upgrade operation finishes, you must unmount the file system cleanly, then re-mount the file system.	Deprecated
Version 7	Version 7 disk layout enables support for variable and large size history log records, more than 2048 volumes, large directory hash, and SmartTier.	Supported
Version 8	Version 8 disk layout enables support for file-level snapshots.	Supported
Version 9	Version 9 disk layout enables support for file compression and, on Solaris SPARC, data deduplication.	Supported

Some of the disk layout versions were not supported on all UNIX operating systems. Currently, only the Version 7, 8, and 9 disk layouts can be created and mounted. The Version 6 disk layout can be mounted, but only for upgrading to a supported version. Disk layout Version 6 cannot be cluster mounted. To cluster mount such a file system, you must first mount the file system on one node and then upgrade to a supported disk layout version using the `vxupgrade` command. No other versions can be created or mounted. Version 9 is the default disk layout version.

The `vxupgrade` command is provided to upgrade an existing VxFS file system to the Version 7, 8, or 9 layout while the file system remains online.

See the `vxupgrade(1M)` manual page.

The `vxfsconvert` command is provided to upgrade Version 1, 2, 3, 4, and 5 disk layouts to the Version 9 disk layout while the file system is not mounted.

See the `vxfsconvert(1M)` manual page.

VxFS Version 7 disk layout

VxFS disk layout Version 7 is similar to Version 6, except that Version 7 enables support for variable and large size history log records, more than 2048 volumes, large directory hash, and SmartTier. The Version 7 disk layout can theoretically support files and file systems up to 8 exabytes (2^{63}). The maximum file system size that can be created is currently restricted to 2^{35} blocks. For a file system to take advantage of greater than 1 terabyte support, it must be created on a Veritas Volume Manager volume. For 64-bit kernels, the maximum size of the file system you can create depends on the block size:

VxFS disk layout Version 7 is similar to Version 6, except that Version 7 enables support for variable and large size history log records, more than 2048 volumes, large directory hash, and SmartTier. The Version 7 disk layout can theoretically support files and file systems up to 8 exabytes (2^{63}). The maximum file system size that can be created is currently restricted to 2^{35} blocks. For a file system to take advantage of greater than 1 terabyte support, it must be created on a Veritas Volume Manager volume. For 64-bit kernels, the maximum size of the file system you can create depends on the block size:

Block Size	Currently-Supported Theoretical Maximum File System Size
1024 bytes	68,719,472,624 sectors (≈32 TB)
2048 bytes	137,438,945,248 sectors (≈64 TB)
4096 bytes	274,877,890,496 sectors (≈128 TB)
8192 bytes	549,755,780,992 sectors (≈256 TB)

The Version 7 disk layout supports group quotas.

See [“About quota files on Veritas File System”](#) on page 924.

Some UNIX commands may not work correctly on file systems larger than one terabyte.

See [“Using UNIX Commands on File Systems Larger than One TB”](#) on page 943.

VxFS Version 8 disk layout

VxFS disk layout Version 8 is similar to Version 7, except that Version 8 enables support for file-level snapshots. The Version 8 disk layout can theoretically support files and file systems up to 8 exabytes (2^{63}). The maximum file system size that can be created is currently restricted to 2^{35} blocks. For a file system to take advantage of greater than 1 terabyte support, it must be created on a Veritas Volume Manager volume. For 64-bit kernels, the maximum size of the file system you can create depends on the block size:

VxFS disk layout Version 8 is similar to Version 7, except that Version 8 enables support for file-level snapshots. The Version 8 disk layout can theoretically support files and file systems up to 8 exabytes (2^{63}). The maximum file system size that can be created is currently restricted to 2^{35} blocks. For a file system to take advantage of greater than 1 terabyte support, it must be created on a Veritas Volume Manager volume. For 64-bit kernels, the maximum size of the file system you can create depends on the block size:

Block Size	Currently-Supported Theoretical Maximum File System Size
1024 bytes	68,719,472,624 sectors (≈ 32 TB)
2048 bytes	137,438,945,248 sectors (≈ 64 TB)
4096 bytes	274,877,890,496 sectors (≈ 128 TB)
8192 bytes	549,755,780,992 sectors (≈ 256 TB)

The Version 8 disk layout supports group quotas.

See [“About quota files on Veritas File System”](#) on page 924.

Some UNIX commands may not work correctly on file systems larger than one terabyte.

See [“Using UNIX Commands on File Systems Larger than One TB”](#) on page 943.

VxFS Version 9 disk layout

VxFS disk layout Version 9 is similar to Version 8, except that Version 9 enables support for data deduplication on Solaris SPARC and file compression. The Version 9 disk layout can theoretically support files and file systems up to 8 exabytes (2^{63}). The maximum file system size that can be created is currently restricted to 2^{35} blocks. For a file system to take advantage of greater than 1 terabyte support, it must be created on a Veritas Volume Manager volume. For 64-bit kernels, the maximum size of the file system you can create depends on the block size:

Block Size	Currently-Supported Theoretical Maximum File System Size
1024 bytes	68,719,472,624 sectors (≈32 TB)
2048 bytes	137,438,945,248 sectors (≈64 TB)
4096 bytes	274,877,890,496 sectors (≈128 TB)
8192 bytes	549,755,780,992 sectors (≈256 TB)

Some UNIX commands may not work correctly on file systems larger than one terabyte.

See [“Using UNIX Commands on File Systems Larger than One TB”](#) on page 943.

Using UNIX Commands on File Systems Larger than One TB

Some UNIX commands may not work correctly on file systems larger than one terabyte.

The `ustat` command returns an `EOVERFLOW` error for VxFS files systems larger than one terabyte because the variable used to store file system size overflows.

See the `ustat(2)` manual page.

System administration utilities such as backup may not operate correctly if they are not large file aware. A large file is a file that is larger than two gigabytes. Similarly, utilities that operate at the file system level must be large file aware to operate correctly on large file systems. A large file system is a file system that is larger than one terabyte. You can have a large file system without creating the file system with the `mkfs -o largefiles` option.

See the `lfcompile(5)` manual page.

Command reference

This appendix includes the following topics:

- [Veritas Volume Manager command reference](#)
- [CVM commands supported for executing on the slave node](#)
- [Veritas Volume Manager manual pages](#)
- [Veritas File System command summary](#)
- [Veritas File System manual pages](#)

Veritas Volume Manager command reference

Most Veritas Volume Manager (VxVM) commands (excepting daemons, library commands and supporting scripts) are linked to the `/usr/sbin` directory from the `/opt/VRTS/bin` directory. It is recommended that you add the following directories to your PATH environment variable:

- If you are using the Bourne or Korn shell (`sh` or `ksh`), use the commands:

```
$ PATH=$PATH:/usr/sbin:/opt/VRTS/bin:/opt/VRTSvxfs/sbin:\
/opt/VRTSdbed/bin:/opt/VRTSob/bin
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

- If you are using a C shell (`csh` or `tcsh`), use the commands:

```
% set path = ( $path /usr/sbin /opt/VRTSvxfs/sbin \
/opt/VRTSdbed/bin /opt/VRTSob/bin /opt/VRTS/bin )
% setenv MANPATH /usr/share/man:/opt/VRTS/man:$MANPATH
```

VxVM library commands and supporting scripts are located under the `/usr/lib/vxvm` directory hierarchy. You can include these directories in your path if you need to use them on a regular basis.

For detailed information about an individual command, refer to the appropriate manual page in the 1M section.

See [“Veritas Volume Manager manual pages”](#) on page 971.

Commands and scripts that are provided to support other commands and scripts, and which are not intended for general use, are not located in `/opt/VRTS/bin` and do not have manual pages.

Commonly-used commands are summarized in the following tables:

- [Table B-1](#) lists commands for obtaining information about objects in VxVM.
- [Table B-2](#) lists commands for administering disks.
- [Table B-3](#) lists commands for creating and administering disk groups.
- [Table B-4](#) lists commands for creating and administering subdisks.
- [Table B-5](#) lists commands for creating and administering plexes.
- [Table B-6](#) lists commands for creating volumes.
- [Table B-7](#) lists commands for administering volumes.
- [Table B-8](#) lists commands for monitoring and controlling tasks in VxVM.

Table B-1 Obtaining information about objects in VxVM

Command	Description
<code>vxctl license [init]</code>	List licensed features of VxVM. The <code>init</code> parameter is required when a license has been added or removed from the host for the new license to take effect.
<code>vxdisk [-g <i>diskgroup</i>] list [<i>diskname</i>]</code>	Lists disks under control of VxVM. See “Displaying disk information” on page 325. Example: <pre># vxdisk -g mydg list</pre>

Table B-1 Obtaining information about objects in VxVM (*continued*)

Command	Description
<code>vx dg list [diskgroup]</code>	Lists information about disk groups. Example: # vx dg list mydg
<code>vx dg -s list</code>	Lists information about shared disk groups. See "Listing shared disk groups" on page 380. Example: # vx dg -s list
<code>vx disk -o alldgs list</code>	Lists all diskgroups on the disks. The imported diskgroups are shown as standard, and additionally all other diskgroups are listed in single quotes.
<code>vx info [-g diskgroup] [volume ...]</code>	Displays information about the accessibility and usability of volumes. See the <i>Veritas Storage Foundation and High Availability Troubleshooting Guide</i> . Example: # vx info -g mydg myvol1 \ myvol2
<code>vx print -hrt [-g diskgroup] [object ...]</code>	Prints single-line information about objects in VxVM. Example: # vx print -g mydg myvol1 \ myvol2
<code>vx list</code>	Provides a consolidated view of the SF configuration, including information from Veritas Volume Manager (VxVM) and Veritas File System (VxFS). See <code>vxlist(1m)</code> manual page.

Table B-1 Obtaining information about objects in VxVM (*continued*)

Command	Description
<code>vxprint -st [-g <i>diskgroup</i>] [<i>subdisk</i> ...]</code>	Displays information about subdisks. Example: # vxprint -st -g mydg
<code>vxprint -pt [-g <i>diskgroup</i>] [<i>plex</i> ...]</code>	Displays information about plexes. Example: # vxprint -pt -g mydg

Table B-2 Administering disks

Command	Description
<code>vxdisk [-o full] reclaim {<i>disk</i> <i>enclosure</i> <i>diskgroup</i>}...</code>	Performs storage reclamation on thin provision LUNs.
<code>vxdiskadm</code>	Administers disks in VxVM using a menu-based interface.
<code>vxdiskadd [<i>devicename</i> ...]</code>	Adds a disk specified by device name. See “Using vxdiskadd to put a disk under VxVM control” on page 345. Example: # vxdiskadd c0t1d0
<code>vxedit [-g <i>diskgroup</i>] rename \ <i>olddisk newdisk</i></code>	Renames a disk under control of VxVM. See “Renaming a disk” on page 348. Example: # vxedit -g mydg rename \ mydg03 mydg02

Table B-2 Administering disks (*continued*)

Command	Description
<code>vxedit [-g <i>diskgroup</i>] set \ reserve=on off <i>diskname</i></code>	Sets aside/does not set aside a disk from use in a disk group. Examples: <pre># vxedit -g mydg set \ reserve=on mydg02 # vxedit -g mydg set \ reserve=off mydg02</pre>
<code>vxedit [-g <i>diskgroup</i>] set \ nohotuse=on off <i>diskname</i></code>	Does not/does allow free space on a disk to be used for hot-relocation. See “Excluding a disk from hot-relocation use” on page 781. See “Making a disk available for hot-relocation use” on page 782. Examples: <pre># vxedit -g mydg set \ nohotuse=on mydg03 # vxedit -g mydg set \ nohotuse=off mydg03</pre>
<code>vxedit [-g <i>diskgroup</i>] set \ spare=on off <i>diskname</i></code>	Adds/removes a disk from the pool of hot-relocation spares. See “Marking a disk as a hot-relocation spare” on page 779. See “Removing a disk from use as a hot-relocation spare” on page 780. Examples: <pre># vxedit -g mydg set \ spare=on mydg04 # vxedit -g mydg set \ spare=off mydg04</pre>
<code>vxdisk offline <i>devicename</i></code>	Takes a disk offline. Example: <pre># vxdisk offline c0t1d0</pre>

Table B-2 Administering disks (*continued*)

Command	Description
<code>vxdg -g <i>diskgroup</i> rmdisk <i>diskname</i></code>	Removes a disk from its disk group. Example: # vxdg -g mydg rmdisk mydg02
<code>vxdiskunsetup <i>devicename</i></code>	Removes a disk from control of VxVM. Example: # vxdiskunsetup c0t3d0

Table B-3 Creating and administering disk groups

Command	Description
<code>vxdg [-s] init <i>diskgroup</i> \</code> <code>[<i>diskname</i>=]<i>devicename</i></code>	Creates a disk group using a pre-initialized disk. See “Creating a shared disk group” on page 381. Example: # vxdg init mydg \ mydg01=c0t1d0
<code>vxdg -g <i>diskgroup</i> listssbinfo</code>	Reports conflicting configuration information. Example: # vxdg -g mydg listssbinfo
<code>vxdg [-n <i>newname</i>] deport <i>diskgroup</i></code>	Deports a disk group and optionally renames it. Example: # vxdg -n newdg deport mydg
<code>vxdg [-n <i>newname</i>] import <i>diskgroup</i></code>	Imports a disk group and optionally renames it. Example: # vxdg -n newdg import mydg

Table B-3 Creating and administering disk groups (*continued*)

Command	Description
<code>vxdg [-n newname] -s import <i>diskgroup</i></code>	Imports a disk group as shared by a cluster, and optionally renames it. See “Importing disk groups as shared” on page 382. Example: <pre># vxdg -n newsdg -s import \ mysdg</pre>
<code>vxdg [-o expand] listmove <i>sourcedg</i> \ <i>targetdg</i> <i>object</i> ...</code>	Lists the objects potentially affected by moving a disk group. See “Listing objects potentially affected by a move” on page 819. Example: <pre># vxdg -o expand listmove \ mydg newdg myvol1</pre>
<code>vxdg [-o expand] move <i>sourcedg</i> \ <i>targetdg</i> <i>object</i> ...</code>	Moves objects between disk groups. See “Moving objects between disk groups” on page 821. Example: <pre># vxdg -o expand move mydg \ newdg myvol1</pre>
<code>vxdg [-o expand] split <i>sourcedg</i> \ <i>targetdg</i> <i>object</i> ...</code>	Splits a disk group and moves the specified objects into the target disk group. See “Splitting disk groups” on page 824. Example: <pre># vxdg -o expand split mydg \ newdg myvol2 myvol3</pre>

Table B-3 Creating and administering disk groups (*continued*)

Command	Description
<code>vx dg join sourcedg targetdg</code>	Joins two disk groups. See “Joining disk groups” on page 825. Example: # <code>vx dg join newdg mydg</code>
<code>vx dg -g diskgroup set \</code> <code>activation=ew ro sr sw off</code>	Sets the activation mode of a shared disk group in a cluster. See “Changing the activation mode on a shared disk group” on page 384. Example: # <code>vx dg -g mysdg set \</code> <code>activation=sw</code>
<code>vx recover -g diskgroup -sb</code>	Starts all volumes in an imported disk group. Example: # <code>vx recover -g mydg -sb</code>
<code>vx dg destroy diskgroup</code>	Destroys a disk group and releases its disks. Example: # <code>vx dg destroy mydg</code>

Table B-4 Creating and administering subdisks

Command	Description
<code>vx make [-g diskgroup] sd subdisk \</code> <code>diskname,offset,length</code>	Creates a subdisk. Example: # <code>vx make -g mydg sd \</code> <code>mydg02-01 mydg02,0,8000</code>

Table B-4 Creating and administering subdisks (*continued*)

Command	Description
<code>vxsd [-g diskgroup] assoc plex \ subdisk...</code>	Associates subdisks with an existing plex. Example: # <code>vxsd -g mydg assoc home-1 \ mydg02-01 mydg02-00 \ mydg02-01</code>
<code>vxsd [-g diskgroup] assoc plex \ subdisk1:0 ... subdiskM:N-1</code>	Adds subdisks to the ends of the columns in a striped or RAID-5 volume. Example: # <code>vxsd -g mydg assoc \ vol01-01 mydg10-01:0 \ mydg11-01:1 mydg12-01:2</code>
<code>vxsd [-g diskgroup] mv oldsubdisk \ newsdisk ...</code>	Replaces a subdisk. Example: # <code>vxsd -g mydg mv mydg01-01 \ mydg02-01</code>
<code>vxsd [-g diskgroup] -s size split \ subdisk sd1 sd2</code>	Splits a subdisk in two. Example: # <code>vxsd -g mydg -s 1000m \ split mydg03-02 mydg03-02 \ mydg03-03</code>
<code>vxsd [-g diskgroup] join \ sd1 sd2 ... subdisk</code>	Joins two or more subdisks. Example: # <code>vxsd -g mydg join \ mydg03-02 mydg03-03 \ mydg03-02</code>

Table B-4 Creating and administering subdisks (*continued*)

Command	Description
<code>vxassist [-g <i>diskgroup</i>] move \</code> <code>volume \!olddisk newdisk</code>	Relocates subdisks in a volume between disks. Example: # <code>vxassist -g mydg move \</code> <code>myvol \!mydg02 mydg05</code> Note: The ! character is a special character in some shells. This example shows how to escape it in a bash shell.
<code>vxunreloc [-g <i>diskgroup</i>] original_disk</code>	Relocates subdisks to their original disks. See “Moving relocated subdisks using vxunreloc” on page 783. Example: # <code>vxunreloc -g mydg mydg01</code>
<code>vxsd [-g <i>diskgroup</i>] dis subdisk</code>	Dissociates a subdisk from a plex. Example: # <code>vxsd -g mydg dis mydg02-01</code>
<code>vxedit [-g <i>diskgroup</i>] rm subdisk</code>	Removes a subdisk. Example: # <code>vxedit -g mydg rm mydg02-01</code>
<code>vxsd [-g <i>diskgroup</i>] -o rm dis subdisk</code>	Dissociates and removes a subdisk from a plex. Example: # <code>vxsd -g mydg -o rm dis \</code> <code>mydg02-01</code>

Table B-5 Creating and administering plexes

Command	Description
<code>vxmake [-g diskgroup] plex plex \ sd=subdisk1[,subdisk2,...]</code>	Creates a concatenated plex. Example: # vxmake -g mydg plex \ vol101-02 \ sd=mydg02-01,mydg02-02
<code>vxmake [-g diskgroup] plex plex \ layout=stripe raid5 stwidth=W \ ncolumn=N \ sd=subdisk1[,subdisk2,...]</code>	Creates a striped or RAID-5 plex. Example: # vxmake -g mydg plex pl-01 \ layout=stripe stwidth=32 \ ncolumn=2 \ sd=mydg01-01,mydg02-01
<code>vxplex [-g diskgroup] att volume plex</code>	Attaches a plex to an existing volume. See “Reattaching a plex manually” on page 852. Example: # vxplex -g mydg att vol101 \ vol101-02
<code>vxplex [-g diskgroup] det plex</code>	Detaches a plex. Example: # vxplex -g mydg det vol101-02
<code>vxmend [-g diskgroup] off plex</code>	Takes a plex offline for maintenance. Example: # vxmend -g mydg off vol102-02

Table B-5 Creating and administering plexes (*continued*)

Command	Description
<code>vxmend [-g diskgroup] on plex</code>	Re-enables a plex for use. See “Reattaching a plex manually” on page 852. Example: <pre># vxmend -g mydg on vol102-02</pre>
<code>vxplex [-g diskgroup] mv oldplex \ newplex</code>	Replaces a plex. Example: <pre># vxplex -g mydg mv \ vol102-02 vol102-03</pre>
<code>vxplex [-g diskgroup] cp volume \ newplex</code>	Copies a volume onto a plex. Example: <pre># vxplex -g mydg cp vol102 \ vol103-01</pre>
<code>vxmend [-g diskgroup] fix clean plex</code>	Sets the state of a plex in an unstartable volume to CLEAN. See “Reattaching a plex manually” on page 852. Example: <pre># vxmend -g mydg fix clean \ vol102-02</pre>
<code>vxplex [-g diskgroup] -o rm dis plex</code>	Dissociates and removes a plex from a volume. Example: <pre># vxplex -g mydg -o rm dis \ vol103-01</pre>

Table B-6 Creating volumes

Command	Description
<code>vxassist [-g <i>diskgroup</i>] maxsize \ layout=<i>layout</i> [<i>attributes</i>]</code>	Displays the maximum size of volume that can be created. Example: # <code>vxassist -g mydg maxsize \ layout=raid5 nlog=2</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume <i>length</i> [layout=<i>layout</i>] \ [<i>attributes</i>]</code>	Creates a volume. See “Creating a volume on specific disks” on page 204. Example: # <code>vxassist -b -g mydg make \ myvol 20g layout=concat \ mydg01 mydg02</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume <i>length</i> layout=mirror \ [nmirror=<i>N</i>] [<i>attributes</i>]</code>	Creates a mirrored volume. See “Creating a mirrored volume” on page 199. Example: # <code>vxassist -b -g mydg make \ mymvol 20g layout=mirror \ nmirror=2</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume <i>length</i> layout=<i>layout</i> \ exclusive=on [<i>attributes</i>]</code>	Creates a volume that may be opened exclusively by a single node in a cluster. See “Creating volumes with exclusive open access by a node” on page 385. Example: # <code>vxassist -b -g mysdg make \ mysmvol 20g layout=mirror \ exclusive=on</code>

Table B-6 Creating volumes (*continued*)

Command	Description
<code>vxassist -b [-g <i>diskgroup</i>] make \</code> <code>volume length layout={stripe raid5} \</code> <code>[stripeunit=<i>W</i>] [ncol=<i>N</i>] \</code> <code>[attributes]</code>	Creates a striped or RAID-5 volume. See “Creating a striped volume” on page 201. See “Creating a RAID-5 volume” on page 203. Example: # <code>vxassist -b -g mydg make \</code> <code>mysvol 20g layout=stripe \</code> <code>stripeunit=32 ncol=4</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \</code> <code>volume length layout=mirror \</code> <code>mirror=ctlr [attributes]</code>	Creates a volume with mirrored data plexes on separate controllers. Example: # <code>vxassist -b -g mydg make \</code> <code>mymcvol 20g layout=mirror \</code> <code>mirror=ctlr</code>
<code>vxmake -b [-g <i>diskgroup</i>] \</code> <code>-Uusage_type vol volume \</code> <code>[len=length] plex=plex,...</code>	Creates a volume from existing plexes. Example: # <code>vxmake -g mydg -Uraid5 \</code> <code>vol r5vol \</code> <code>plex=raidplex,raidlog1,\</code> <code>raidlog2</code>
<code>vxvol [-g <i>diskgroup</i>] start volume</code>	Initializes and starts a volume for use. Example: # <code>vxvol -g mydg start r5vol</code>
<code>vxvol [-g</code> <code>diskgroup] init zero \</code> <code>volume</code>	Initializes and zeros out a volume for use. Example: # <code>vxvol -g mydg init zero \</code> <code>myvol</code>

Table B-7 Administering volumes

Command	Description
<code>vxassist [-g <i>diskgroup</i>] mirror \</code> <code>volume [<i>attributes</i>]</code>	Adds a mirror to a volume. See “Adding a mirror to a volume” on page 838. Example: <pre># vxassist -g mydg mirror \ myvol mydg10</pre>
<code>vxassist [-g <i>diskgroup</i>] remove \</code> <code>mirror volume [<i>attributes</i>]</code>	Removes a mirror from a volume. See “Removing a mirror” on page 842. Example: <pre># vxassist -g mydg remove \ mirror myvol \!mydg11</pre> Note: The ! character is a special character in some shells. This example shows how to escape it in a bash shell.
<code>vxassist [-g <i>diskgroup</i>] \</code> <code>{growto growby} volume length</code>	Grows a volume to a specified size or by a specified amount. Example: <pre># vxassist -g mydg growby \ myvol 10g</pre>
<code>vxassist [-g <i>diskgroup</i>] \</code> <code>{shrinkto shrinkby} volume length</code>	Shrinks a volume to a specified size or by a specified amount. Example: <pre># vxassist -g mydg shrinkto \ myvol 20g</pre>

Table B-7 Administering volumes (*continued*)

Command	Description
<code>vxresize -b -F vxfs [-g diskgroup] \ volume length diskname ...</code>	Resizes a volume and the underlying Veritas File System. Example: # <code>vxresize -b -F vxfs \ -g mydg myvol 20g mydg10 \ mydg11</code>
<code>vxsnap [-g diskgroup] prepare volume [drl=on sequential off]</code>	Prepares a volume for instant snapshots and for DRL logging. See “Adding an instant snap DCO and DCO volume” on page 596. Example: # <code>vxsnap -g mydg prepare \ myvol drl=on</code>
<code>vxsnap [-g diskgroup] make \ source=volume\ /newvol=snapvol\ [/nmirror=number]</code>	Takes a full-sized instant snapshot of a volume by breaking off plexes of the original volume. See “Creating instant snapshots” on page 595. Example: # <code>vxsnap -g mydg make \ source=myvol/ newvol=mysnpvol/ nmirror=2</code>

Table B-7 Administering volumes (*continued*)

Command	Description
<code>vxsnap [-g diskgroup] make \</code> <code>source=volume/snapvol=snapvol</code>	Takes a full-sized instant snapshot of a volume using a prepared empty volume. See “Creating a volume for use as a full-sized instant or linked break-off snapshot” on page 600. See “Creating instant snapshots” on page 595. Example: <pre># vxsnap -g mydg make \ source=myvol/snapvol=snpvol</pre>
<code>vxmake [-g diskgroup] cache \</code> <code>cache_object cachevolname=volume \</code> <code>[regionsize=size]</code>	Creates a cache object for use by space-optimized instant snapshots. See “Creating a shared cache object” on page 598. A cache volume must have already been created. After creating the cache object, enable the cache object with the <code>vxcache start</code> command. For example: <pre># vxassist -g mydg make \ cvol 1g layout=mirror \ init=active mydg16 mydg17 # vxmake -g mydg cache cobj \ cachevolname=cvol # vxcache -g mydg start cobj</pre>
<code>vxsnap [-g diskgroup] make \</code> <code>source=volume/newvol=snapvol\</code> <code>/cache=cache_object</code>	Takes a space-optimized instant snapshot of a volume. See “Creating instant snapshots” on page 595. Example: <pre># vxsnap -g mydg make \ source=myvol/\ newvol=mysosvol/\ cache=cobj</pre>

Table B-7 Administering volumes (*continued*)

Command	Description
<code>vxsnap [-g <i>diskgroup</i>] refresh <i>snapshot</i></code>	Refreshes a snapshot from its original volume. See “Refreshing an instant space-optimized snapshot” on page 616. Example: <pre># vxsnap -g mydg refresh \ mysnpvol</pre>
<code>vxsnap [-g <i>diskgroup</i>] dis <i>snapshot</i></code>	Turns a snapshot into an independent volume. See “Dissociating an instant snapshot” on page 618. Example: <pre># vxsnap -g mydg dis mysnpvol</pre>
<code>vxsnap [-g <i>diskgroup</i>] unprepare \ <i>volume</i></code>	Removes support for instant snapshots and DRL logging from a volume. Example: <pre># vxsnap -g mydg unprepare \ myvol</pre>
<code>vxassist [-g <i>diskgroup</i>] relayout \ <i>volume</i> [<i>layout=layout</i>] \ [<i>relayout_options</i>]</code>	Performs online relayout of a volume. See “Performing online relayout” on page 831. Example: <pre># vxassist -g mydg relayout \ vol2 layout=stripe</pre>

Table B-7 Administering volumes (*continued*)

Command	Description
<pre>vxassist [-g <i>diskgroup</i>] relayout \ volume layout=raid5 \ stripeunit=<i>W</i> \ ncol=<i>N</i></pre>	Relays out a volume as a RAID-5 volume with stripe width <i>W</i> and <i>N</i> columns. See “Performing online relayout” on page 831. Example: <pre># vxassist -g mydg relayout \ vol3 layout=raid5 \ stripeunit=16 ncol=4</pre>
<pre>vxrelayout [-g <i>diskgroup</i>] -o bg \ reverse volume</pre>	Reverses the direction of a paused volume relayout. See “Volume sets” on page 142. Example: <pre># vxrelayout -g mydg -o bg \ reverse vol3</pre>
<pre>vxassist [-g <i>diskgroup</i>] convert \ volume [layout=<i>layout</i>] \ [convert_options]</pre>	Converts between a layered volume and a non-layered volume layout. Example: <pre># vxassist -g mydg convert \ vol3 layout=stripe-mirror</pre>
<pre>vxassist [-g <i>diskgroup</i>] remove \ volume volume</pre>	Removes a volume. See “Removing a volume” on page 854. Example: <pre># vxassist -g mydg remove \ myvol</pre>

Table B-8 Monitoring and controlling tasks

Command	Description
<code>command [-g diskgroup] -t tasktag \</code> <code>[options] [arguments]</code>	Specifies a task tag to a VxVM command. See “Specifying task tags” on page 827. Example: <pre># vxrecover -g mydg \ -t mytask -b mydg05</pre>
<code>vxtask [-h] [-g diskgroup] list</code>	Lists tasks running on a system. See “Using the vxtask command” on page 830. Example: <pre># vxtask -h -g mydg list</pre>
<code>vxtask monitor task</code>	Monitors the progress of a task. See “Using the vxtask command” on page 830. Example: <pre># vxtask monitor mytask</pre>
<code>vxtask pause task</code>	Suspends operation of a task. See “Using the vxtask command” on page 830. Example: <pre># vxtask pause mytask</pre>
<code>vxtask -p [-g diskgroup] list</code>	Lists all paused tasks. See “Using the vxtask command” on page 830. Example: <pre># vxtask -p -g mydg list</pre>

Table B-8 Monitoring and controlling tasks *(continued)*

Command	Description
<code>vxtask resume task</code>	Resumes a paused task. See “Using the vxtask command” on page 830. Example: <pre># vxtask resume mytask</pre>
<code>vxtask abort task</code>	Cancels a task and attempts to reverse its effects. See “Using the vxtask command” on page 830. Example: <pre># vxtask abort mytask</pre>

CVM commands supported for executing on the slave node

[Table B-9](#) shows the complete list of commands that are supported for executing on the slave node.

Table B-9 List of CVM commands supported for executing on the slave node

Command	Supported operations
vxdg	

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
	vxdg -s init <shared_dg> [cds=on off] vxdg -T <different_versions> -s init <shared_dg> [minor=base-minor] [cds=on off] vxdg [-n newname] [-h new-host-id] deport <shared_dg> vxdg [-Cfst] [-n newname] [-o clearreserve] [-o useclonedev={on off}] [-o updateid] [-o noreonline] [-o selectcp=diskid] [-o dgtype=shared] import <shared_dg> vxdg destroy <shared_dg> vxdg -g <shared_dg> [-o overridesb] [-f] adddisk [disk=]device vxdg -g <shared_dg> addsite site vxdg -g <shared_dg> reattachsite site vxdg -g <shared_dg> detachsite site vxdg -g <shared_dg> rmsite site vxdg -g <shared_dg> renamesite oldname newname vxdg flush <shared_dg> vxdg [-qa] -g <shared_dg> free [medianame...] vxdg join sourcedg targetdg (both dgs should be shared) vxdg split sourcedg targetdg vxdg [-q] [-s] [-o listreserve] list [diskgroup...] vxdg [-o expand] move sourcedg targetdg object (both dgs should be shared) vxdg -g shared_dg recover vxdg -g <shared_dg> [-f] remenor <shared_dg> new-minor-number vxdg -g <shared_dg> rmdisk medianame... vxdg -g <shared_dg> [-q] spare [medianame...] vxdg -g <shared_dg> [-f] [-o retain replace] settag [encl:<enclosure>] name[=value name[=value] vxdg [-q] listtag <shared_dg> vxdg -g <shared_dg> rmtag [encl:<enclosure>] name=value vxdg -g <shared_dg> set siteconsistent=on vxdg upgrade <shared_dg>

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
	vxdbg -g <shared_dg>set attr=value ...
vxassist	vxassist -g <shared_dg> [-b] convert volume layout=<type> vxassist -g <shared_dg> [-b] addlog volume vxassist -g <shared_dg> [-b] mirror volume vxassist [-b]-g <shared_dg>make volume length [layout=layout] diskname ... vxassist -g <shared_dg> [-b] growby volume lengthchange [attribute ...] vxassist [-b] -g <shared_dg> growto volume newlength vxassist -g <shared_dg> shrinkby volume lengthchange vxassist -g <shared_dg> shrinkto volume newlength vxassist -g <shared_dg> settag volume vset tagname[=tagvalue] vxassist -g <shared_dg>replacetag volume vset oldtag newtag vxassist -g <shared_dg> removetag volume vset tagname vxassist -g <shared_dg> move volume-name storage-spec vxassist -g <shared_dg> relayout {volume-name} layout=<type> vxassist -g <shared_dg> remove {volume mirror log} volume-name vxassist -g <shared_dg> snapshot volume-name [snapshot-name] [comment=<comment>] vxassist -g <shared_dg> snapstart volume vxassist -g <shared_dg> maxsize layout=<> nmirror=<> / nlog=<> vxassist -g <shared_dg> maxgrow volume vxassist -g <shared_dg> snapback snapvol vxassist -g <shared_dg> snapclear snapvol1

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
vxcache	vxcache -g <shared_dg> start cacheobject vxcache -g <shared_dg> stop cacheobject vxcache -g <shared_dg> att volume cacheobject vxcache -g <shared_dg> dis cachevol vxcache -g <shared_dg> shrinkcacheto cacheobject newlength vxcache -g <shared_dg> shrinkcacheby cacheobject lengthchange vxcache -g <shared_dg> growcacheto cacheobject newlength vxcache -g <shared_dg> growcacheby cacheobject lengthchange
vxdco	vxdco -g <shared_dg> dis dco vxdco -g <shared_dg> att volume dco vxdco -g <shared_dg>[-o force] enable dco
vxedit	vxedit -g <shared_dg> set comment="plex comment" plex1 vxedit -g <shared_dg> -rf rm volume vxedit -g <shared_dg>rename oldname newname vxedit -g <shared_dg> set what=value vxedit-g <shared_dg> set user=value mode=value medianame vxedit -g <shared_dg> set failing=off <disk name> vxedit -g <shared_dg> set fstype volumename vxedit -g <shared_dg> set len subdisk vxedit -g <shared_dg> set orig_dmname subdisk vxedit -g <shared_dg> set orig_dmooffset subdisk vxedit -g <shared_dg> set diskdetpolicy diskgroup
vxmake	vxmake -g <shared_dg> sd name [attr...] vxmake-g <shared_dg>plex plex sd=subdisk1[,subdisk2,...] vxmake -g <shared_dg> -U fsgen vol homevol1 plex=plex-1 vxmake -g <shared_dg> -U fsgen vol volume1 plex=plex1,plex2 vxmake -g <shared_dg> cache name regionsize=<size> vxmake -g <shared_dg> dco volume log=dco

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
vxmend	vxmend -g <shared_dg> on plex vxmend -g <shared_dg> off plex
vxmirror	vxmirror -g <shared_dg> medianame vxmirror -g <shared_dg> -d [yes no]
vxplex	vxplex -g <shared_dg> att volume plex vxplex -g <shared_dg> cp volume new_plex vxplex -g <shared_dg> dis plex1 vxplex -g <shared_dg> mv original_plex new_plex vxplex -g <shared_dg> snapstart vol snapplex vxplex -g <shared_dg> snapshot snapplex vxplex -g <shared_dg> snapback vol snapplex vxplex -g <shared_dg> plex
vxrelayout	vxrelayout -g <shared_dg> status volume vxrelayout -g <shared_dg> start volume vxrelayout -g <shared_dg> reverse volname
vxsd	vxsd -g <shared_dg> assoc plex subdisk1 [subdisk2 subdisk3 ...] vxsd -g <shared_dg> [-o force] dis subdisk vxsd -g <shared_dg> mv old_subdisk new_subdisk [new_subdisk ...] vxsd -g <shared_dg> aslog plex2 sdisk3 vxsd -g <shared_dg> join subdisk1 subdisk2 ... new_subdisk vxsd -g <shared_dg> [-o force] dis subdisk vxsd -g <shared_dg> split subdisk newsd [newsd2]...

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
vxsnap	vxsnap -g <shared_dg> addmir volume [nmirror=N] vxsnap -g <shared_dg> prepare volume vxsnap -g <shared_dg> rmmir volume vxsnap -g <shared_dg> unprepare volume vxsnap -g <shared_dg> make snapshot_tuple [snapshot_tuple]... [alloc=storage_attributes] vxsnap [-f] -g <shared_dg> dis volume vxsnap -g <shared_dg> addmap volumenam count vxsnap -g <shared_dg> print volumenam vxsnap -g <shared_dg> list volumenam vxsnap -g <shared_dg> syncwait snapvol vxsnap -g <shared_dg> snapwait vxsnap -g <shared_dg> refresh snapvol source=volume vxsnap -g <shared_dg> restore target source=volname vxsnap -g <shared_dg> split volumenam vxsnap -g <shared_dg> reattach volname source=volname
vxsnptadm	vxsnptadm -g <shared_dg> create vol [snptname=snpt] [snapvolname=snapvol] [data={yes no}] vxsnptadm -g <shared_dg> info vol [snptname=snpt] vxsnptadm -g <shared_dg> remove vol snptname=snpt vxsnptadm -g <shared_dg> removeall vol [cookie=cookie] vxsnptadm -g <shared_dg> rename vol snptname=snpt newname=snpt2

Table B-9 List of CVM commands supported for executing on the slave node
(continued)

Command	Supported operations
vxvol	vxvol -g <shared_dg> set logtype=drl drlseq volume vxvol -g <shared_dg> start volume vxvol -g <shared_dg> stop volume vxvol -g <shared_dg> {startall stopall} volume vxvol -g <shared_dg> init enable volume vxvol -g <shared_dg> init active volume vxvol -g <shared_dg> maint volumename vxvol -g <shared_dg> set len volumename vxvol -g <shared_dg> set logtype volumename vxvol -g <shared_dg> set loglen volumename
vxvset	vxvset -g <shared_dg> make volume-set-name volume-name vxvset -g <shared_dg> addvol volume-set-name volume-name vxvset -g <shared_dg> list volume-set-name vxvset -g <shared_dg> rmvol volume-set-name volume-name vxvset -g <shared_dg> stop volume-set-name vxvset -g <shared_dg> start volume-set-name
vxevac	vxevac -g <shared_dg> medianame
vxresize	vxresize [-Vsb] [-F fstype] -g <shared_dg> volume length
vxrecover	vxrecover -g <shared_dg> vxrecover -g <shared_dg> volume
vxckdiskrm	vxckdiskrm -g <shared_dg> medianame

Veritas Volume Manager manual pages

Manual pages are organized into the following sections:

- 1M Administrative commands.
- 4 File formats.

7 Device driver interfaces.

Section 1M — administrative commands

[Table B-10](#) lists the manual pages in section 1M for commands that are used to administer Veritas Volume Manager.

Table B-10 Section 1M manual pages

Name	Description
<code>vxapslice</code>	Manage an area of disk for use by an Alternate Pathing database.
<code>vxassist</code>	Create, relayout, convert, mirror, backup, grow, shrink, delete, and move volumes.
<code>vxbootsetup</code>	Set up system boot information on a Veritas Volume Manager disk.
<code>vxcache</code>	Administer the cache object for space-optimized snapshots.
<code>vxcached</code>	Resize cache volumes when required.
<code>vxcdsconvert</code>	Make disks and disk groups portable between systems.
<code>vxclustadm</code>	Start, stop, and reconfigure a cluster.
<code>vxcmdlog</code>	Administer command logging.
<code>vxconfigbackup</code>	Back up disk group configuration.
<code>vxconfigbackupd</code>	Disk group configuration backup daemon.
<code>vxconfigd</code>	Veritas Volume Manager configuration daemon
<code>vxconfigrestore</code>	Restore disk group configuration.
<code>vxdarestore</code>	Restore simple or nopriv disk access records.
<code>vxdco</code>	Perform operations on version 0 DCO objects and DCO volumes.
<code>vxctl</code>	Control the volume configuration daemon.

Table B-10 Section 1M manual pages (*continued*)

Name	Description
<code>vxddladm</code>	Device Discovery Layer subsystem administration.
<code>vxdefault</code>	Manage the defaults set in <code>/etc/default/vxsf</code> that configure settings such as SmartMove, thin reclamation, automatic starting of volumes, and minor numbers for shared disk groups.
<code>vxdg</code>	Manage Veritas Volume Manager disk groups.
<code>vxdisk</code>	Define and manage Veritas Volume Manager disks.
<code>vxdiskadd</code>	Add one or more disks for use with Veritas Volume Manager.
<code>vxdiskadm</code>	Menu-driven Veritas Volume Manager disk administration.
<code>vxdiskconfig</code>	Configure disk devices and bring them under VxVM control.
<code>vxdisksetup</code>	Configure a disk for use with Veritas Volume Manager.
<code>vxdiskunsetup</code>	Deconfigure a disk from use with Veritas Volume Manager.
<code>vxdmppadm</code>	DMP subsystem administration.
<code>vxedit</code>	Create, remove, and modify Veritas Volume Manager records.
<code>vxencap</code>	Encapsulate partitions on a new disk.
<code>vxevac</code>	Evacuate all volumes from a disk.
<code>vxinfo</code>	Print accessibility and usability of volumes.
<code>vxinstall</code>	Menu-driven Veritas Volume Manager initial configuration.
<code>vxintro</code>	Introduction to the Veritas Volume Manager utilities.

Table B-10 Section 1M manual pages (*continued*)

Name	Description
<code>vxiod</code>	Start, stop, and report on Veritas Volume Manager kernel I/O threads.
<code>vxlufinish</code>	Finish a live upgrade of Veritas Volume Manager.
<code>vxlustart</code>	Start a live upgrade of Veritas Volume Manager.
<code>vxmake</code>	Create Veritas Volume Manager configuration records.
<code>vxmemstat</code>	Display memory statistics for Veritas Volume Manager.
<code>vxmend</code>	Mend simple problems in configuration records.
<code>vxmirror</code>	Mirror volumes on a disk or control default mirroring.
<code>vxnotify</code>	Display Veritas Volume Manager configuration events.
<code>vxplex</code>	Perform Veritas Volume Manager operations on plexes.
<code>vxprint</code>	Display records from the Veritas Volume Manager configuration.
<code>vxr5check</code>	Verify RAID-5 volume parity.
<code>vxreattach</code>	Reattach disk drives that have become accessible again.
<code>vxrecover</code>	Perform volume recovery operations.
<code>vxrelayout</code>	Convert online storage from one layout to another.
<code>vxrelocd</code>	Monitor Veritas Volume Manager for failure events and relocate failed subdisks.
<code>vxresize</code>	Change the length of a volume containing a file system.
<code>vxrootadm</code>	Grow or take snapshots of the boot disk.

Table B-10 Section 1M manual pages (*continued*)

Name	Description
<code>vxscsiinq</code>	Display SCSI inquiry data.
<code>vxsd</code>	Perform Veritas Volume Manager operations on subdisks.
<code>vxsnap</code>	Enable DRL on a volume, and create and administer instant snapshots.
<code>vxsparecheck</code>	Monitor for disk failure, and replace failed disks.
<code>vxstat</code>	Veritas Volume Manager statistics management utility.
<code>vxtask</code>	List and administer Veritas Volume Manager tasks.
<code>vxtrace</code>	Trace operations on volumes.
<code>vxtranslog</code>	Administer transaction logging.
<code>vxtune</code>	Adjust Veritas Volume Replicator and Veritas Volume Manager tunables.
<code>vxunreloc</code>	Move a hot-relocated subdisk back to its original disk.
<code>vxunroot</code>	Remove Veritas Volume Manager hooks from encapsulated root volumes.
<code>vxvol</code>	Perform Veritas Volume Manager operations on volumes.
<code>vxvset</code>	Create and administer volume sets.

Section 4 — file formats

[Table B-11](#) lists the manual pages in section 4 that describe the format of files that are used by Veritas Volume Manager.

Table B-11 Section 4 manual pages

Name	Description
<code>vol_pattern</code>	Disk group search specifications.

Table B-11 Section 4 manual pages (*continued*)

Name	Description
<code>vxmake</code>	<code>vxmake</code> description file.

Section 7 — device driver interfaces

[Table B-12](#) lists the manual pages in section 7 that describe the interfaces to Veritas Volume Manager devices.

Table B-12 Section 7 manual pages

Name	Description
<code>vxconfig</code>	Configuration device.
<code>vxdmp</code>	Dynamic Multi-Pathing device.
<code>vxinfo</code>	General information device.
<code>vxio</code>	Virtual disk device.
<code>vxiod</code>	I/O daemon process control device.
<code>vxtrace</code>	I/O tracing device.

Veritas File System command summary

Symbolic links to all VxFS command executables are installed in the `/opt/VRTS/bin` directory. Add this directory to the end of your `PATH` environment variable to access the commands.

[Table B-13](#) describes the VxFS-specific commands.

Table B-13 VxFS commands

Command	Description
<code>cp</code>	Copies files and directories on VxFS file systems.
<code>df</code>	Reports the number of free disk blocks and inodes for a VxFS file system.
<code>fcldm</code>	Administers VxFS File Change Logs.
<code>ff</code>	Lists file names and inode information for a VxFS file system.
<code>fiostat</code>	Administers file I/O statistics

Table B-13 VxFS commands (*continued*)

Command	Description
<code>fsadm</code>	Resizes or defragments a VxFS file system.
<code>fsapadm</code>	Administers VxFS allocation policies.
<code>fscat</code>	Cats a VxFS file system.
<code>fscdsadm</code>	Performs online CDS operations.
<code>fscdsconv</code>	Performs offline CDS migration tasks on VxFS file systems.
<code>fscdstask</code>	Performs various CDS operations.
<code>fsck</code>	Checks and repairs a VxFS file system.
<code>fsckpt_restore</code>	Restores file systems from VxFS Storage Checkpoints.
<code>fsclustadm</code>	Manages cluster-mounted VxFS file systems. This functionality is available only with the Veritas Storage Foundation Cluster File System High Availability product.
<code>fsdb</code>	Debugs VxFS file systems.
<code>fsdedupadm</code>	Administers data deduplication. Solaris SPARC, only.
<code>fsfreeze</code>	Freezes VxFS file systems and executes a user command on the file systems.
<code>fsmap</code>	Displays VxFS file system extent information.
<code>fsmigadm</code>	Administers file system online migrations.
<code>fsppadm</code>	Administers VxFS placement policies.
<code>fsppmk</code>	Creates placement policies.
<code>fstag</code>	Creates, deletes, or lists file tags.
<code>fstyp</code>	Returns the type of file system on a specified disk partition.
<code>fsvmap</code>	Maps volumes of VxFS file systems to files.
<code>fsvoladm</code>	Administers VxFS volumes.
<code>getext</code>	Gets extent attributes for a VxFS file system.
<code>glmconfig</code>	Configures Group Lock Managers (GLM). This functionality is available only with the Veritas Storage Foundation Cluster File System High Availability product.

Table B-13 VxFS commands (*continued*)

Command	Description
glmdump	Reports stuck Group Lock Managers (GLM) locks in a cluster file system.
glmstat	Group Lock Managers (GLM) statistics gathering utility. This functionality is available only with the Veritas Storage Foundation Cluster File System High Availability product.
ls	Lists files and directories on a VxFS file system.
mkdstfs	SmartTier file system creation utility.
mkfs	Constructs a VxFS file system.
mount	Mounts a VxFS file system.
mv	Moves files and directories on a VxFS file system.
ncheck	Generates path names from inode numbers for a VxFS file system.
qioadmin	Administers VxFS Quick I/O for Databases cache.
qiomkfile	Creates a VxFS Quick I/O device file. This functionality is available only with the Veritas Quick I/O for Databases feature.
qiostat	Displays statistics for VxFS Quick I/O for Databases. This functionality is available only with the Veritas Quick I/O for Databases feature.
setext	Sets extent attributes on a file in a VxFS file system.
umount	Unmounts a VxFS file system.
vxcompress	Compresses and uncompresses files.
vxdump	Incrementally dumps file systems.
vxedquota	Edits user quotas for a VxFS file system.
vxenablf	Enables specific VxFS features.
vxfilesnap	Makes a copy-on-write copy of a file in a VxFS file system.
vxfsconvert	Converts an unmounted file system to VxFS or upgrades a VxFS disk layout version.
vxfsstat	Displays file system statistics.
vxlsino	Looks up VxFS reverse path names.
vxquot	Displays file system ownership summaries for a VxFS file system.

Table B-13 VxFS commands (*continued*)

Command	Description
<code>vxquota</code>	Displays user disk quotas and usage on a VxFS file system.
<code>vxquotaoff</code> <code>vxquotaon</code>	Turns quotas on and off for a VxFS file system.
<code>vxrepquota</code>	Summarizes quotas for a VxFS file system.
<code>vxrestore</code>	Restores a file system incrementally.
<code>vxtunefs</code>	Tunes a VxFS file system.
<code>vxupgrade</code>	Upgrades the disk layout of a mounted VxFS file system.

Veritas File System manual pages

This release includes the following online manual pages as part of the `VRTSvxfs` package. These are installed in the appropriate directories under `/opt/VRTS/man` (add this to your `MANPATH` environment variable), but does not update the `windex` database. To ensure that new VxFS manual pages display correctly, update the `windex` database after installing `VRTSvxfs`.

See the `catman(1M)` manual page.

[Table B-14](#) describes the VxFS-specific section 1 manual pages.

Table B-14 Section 1 manual pages

Section 1	Description
<code>cp_vxfs</code>	Copies files and directories on a VxFS file system.
<code>cpio_vxfs</code>	Copies files and directories on a VxFS file system.
<code>fiostat</code>	Administers file I/O statistics.
<code>fsmap</code>	Displays VxFS file system extent information.
<code>getext</code>	Gets extent attributes for a VxFS file system.
<code>ls_vxfs</code>	Lists files and directories on a VxFS file system.
<code>mv_vxfs</code>	Moves files and directories on a VxFS file system.
<code>qioadmin</code>	Administers VxFS Quick I/O for Databases cache. This functionality is available only with the Veritas Quick I/O for Databases feature.

Table B-14 Section 1 manual pages (*continued*)

Section 1	Description
qiomkfile	Creates a VxFS Quick I/O device file. This functionality is available only with the Veritas Quick I/O for Databases feature.
qiostat	Displays statistics for VxFS Quick I/O for Databases. This functionality is available only with the Veritas Quick I/O for Databases feature.
setext	Sets extent attributes on a file in a VxFS file system.
vxcompress	Compresses or uncompresses files.
vxfilesnap	Makes a copy-on-write copy of a file in a VxFS file system.

[Table B-15](#) describes the VxFS-specific section 1M manual pages.

Table B-15 Section 1M manual pages

Section 1M	Description
df_vxfs	Reports the number of free disk blocks and inodes for a VxFS file system.
fcladm	Administers VxFS File Change Logs.
ff_vxfs	Lists file names and inode information for a VxFS file system.
fsadm_vxfs	Resizes or reorganizes a VxFS file system.
fsapadm	Administers VxFS allocation policies.
fscat_vxfs	Cats a VxFS file system.
fscdsadm	Performs online CDS operations.
fscdsconv	Performs offline CDS migration tasks on VxFS file systems.
fscdstask	Performs various CDS operations.
fsck_vxfs	Checks and repairs a VxFS file system.
fsckptadm	Performs various administrative tasks like creating, deleting, converting, setting, and displaying the quota on a Storage Checkpoint. Quota display can be formatted in a human-friendly way, using the <code>-H</code> option.
fsckpt_restore	Restores file systems from VxFS Storage Checkpoints.
fsclustadm	Manages cluster-mounted VxFS file systems. This functionality is available only with the Veritas Cluster File System product.

Table B-15 Section 1M manual pages (*continued*)

Section 1M	Description
<code>fsdbencap</code>	Encapsulates databases.
<code>fsdb_vxfs</code>	Debugs VxFS file systems.
<code>fsdedupadm</code>	Administers data deduplication. Solaris SPARC, only.
<code>fsfreeze</code>	Freezes VxFS file systems and executes a user command on the file systems.
<code>fsmigadm</code>	Administers file system online migrations.
<code>fspadm</code>	Administers VxFS placement policies.
<code>fstyp_vxfs</code>	Returns the type of file system on a specified disk partition.
<code>fsvmap</code>	Maps volumes of VxFS file systems to files.
<code>fsvoladm</code>	Administers VxFS volumes.
<code>glmconfig</code>	Configures Group Lock Managers (GLM). This functionality is available only with the Veritas Cluster File System product.
<code>glmdump</code>	Reports stuck Group Lock Managers (GLM) locks in a cluster file system.
<code>mkdstfs</code>	SmartTier file system creation utility.
<code>mkfs_vxfs</code>	Constructs a VxFS file system.
<code>mount_vxfs</code>	Mounts a VxFS file system.
<code>ncheck_vxfs</code>	Generates path names from inode numbers for a VxFS file system.
<code>quot</code>	Summarizes ownership on a VxFS file system.
<code>quotacheck_vxfs</code>	Checks VxFS file system quota consistency.
<code>umount_vxfs</code>	Unmounts a VxFS file system.
<code>vxdiskusg</code>	Generates VxFS disk accounting data by user ID.
<code>vxdump</code>	Incrementally dumps file systems.
<code>vxedquota</code>	Edits user quotas for a VxFS file system.
<code>vxenablef</code>	Enables specific VxFS features.
<code>vxfsconvert</code>	Converts an unmounted file system to VxFS disk layout.

Table B-15 Section 1M manual pages (*continued*)

Section 1M	Description
<code>vxfsconvert</code>	Converts an unmounted file system to VxFS or upgrades a VxFS disk layout version.
<code>vxfsstat</code>	Displays file system statistics.
<code>vxlsino</code>	Looks up VxFS reverse path names.
<code>vxquot</code>	Displays file system ownership summaries for a VxFS file system.
<code>vxquota</code>	Displays user disk quotas and usage on a VxFS file system.
<code>vxquotaoff</code> <code>vxquotaon</code>	Turns quotas on and off for a VxFS file system.
<code>vxrepquota</code>	Summarizes quotas for a VxFS file system.
<code>vxrestore</code>	Restores a file system incrementally.
<code>vxtunefs</code>	Tunes a VxFS file system.
<code>vxupgrade</code>	Upgrades the disk layout of a mounted VxFS file system.

[Table B-16](#) describes the VxFS-specific section 3 manual pages.

Table B-16 Section 3 manual pages

Section 3	Description
<code>vxfs_ap_alloc2</code>	Allocates an <code>fsap_info2</code> structure.
<code>vxfs_ap_assign_ckpt</code>	Assigns an allocation policy to file data and metadata in a Storage Checkpoint.
<code>vxfs_ap_assign_ckptchain</code>	Assigns an allocation policy for all of the Storage Checkpoints of a VxFS file system.
<code>vxfs_ap_assign_ckptdef</code>	Assigns a default allocation policy for new Storage Checkpoints of a VxFS file system.
<code>vxfs_ap_assign_file</code>	Assigns an allocation policy for file data and metadata.
<code>vxfs_ap_assign_file_pat</code>	Assigns a pattern-based allocation policy for a directory.
<code>vxfs_ap_assign_fs</code>	Assigns an allocation policy for all file data and metadata within a specified file system.
<code>vxfs_ap_assign_fs_pat</code>	Assigns an pattern-based allocation policy for a file system.

Table B-16 Section 3 manual pages (*continued*)

Section 3	Description
<code>vxfs_ap_define</code>	Defines a new allocation policy.
<code>vxfs_ap_define2</code>	Defines a new allocation policy.
<code>vxfs_ap_enforce_ckpt</code>	Reorganizes blocks in a Storage Checkpoint to match a specified allocation policy.
<code>vxfs_ap_enforce_ckptchain</code>	Enforces the allocation policy for all of the Storage Checkpoints of a VxFS file system.
<code>vxfs_ap_enforce_file</code>	Ensures that all blocks in a specified file match the file allocation policy.
<code>vxfs_ap_enforce_file2</code>	Reallocates blocks in a file to match allocation policies.
<code>vxfs_ap_enforce_range</code>	Reallocates blocks in a file within a specified range to match allocation policies.
<code>vxfs_ap_enumerate</code>	Returns information about all allocation policies.
<code>vxfs_ap_enumerate2</code>	Returns information about all allocation policies.
<code>vxfs_ap_free2</code>	Frees one or more <code>fsap_info2</code> structures.
<code>vxfs_ap_query</code>	Returns information about a specific allocation policy.
<code>vxfs_ap_query2</code>	Returns information about a specific allocation policy.
<code>vxfs_ap_query_ckpt</code>	Returns information about allocation policies for each Storage Checkpoint.
<code>vxfs_ap_query_ckptdef</code>	Retrieves the default allocation policies for new Storage Checkpoints of a VxFS file system
<code>vxfs_ap_query_file</code>	Returns information about allocation policies assigned to a specified file.
<code>vxfs_ap_query_file_pat</code>	Returns information about the pattern-based allocation policy assigned to a directory.
<code>vxfs_ap_query_fs</code>	Retrieves allocation policies assigned to a specified file system.
<code>vxfs_ap_query_fs_pat</code>	Returns information about the pattern-based allocation policy assigned to a file system.
<code>vxfs_ap_remove</code>	Deletes a specified allocation policy.
<code>vxfs_fcl_sync</code>	Sets a synchronization point in the VxFS File Change Log.
<code>vxfs_fiostats_dump</code>	Returns file and file range I/O statistics.

Table B-16 Section 3 manual pages (*continued*)

Section 3	Description
<code>vxfs_fiostats_getconfig</code>	Gets file range I/O statistics configuration values.
<code>vxfs_fiostats_set</code>	Turns on and off file range I/O statistics and resets statistics counters.
<code>vxfs_get_iooffsets</code>	Obtains VxFS inode field offsets.
<code>vxfs_inotopath</code>	Returns path names for a given inode number.
<code>vxfs_inostat</code>	Gets the file statistics based on the inode number.
<code>vxfs_inotofd</code>	Gets the file descriptor based on the inode number.
<code>vxfs_nattr_check</code> <code>vxfs_nattr_fcheck</code>	Checks for the existence of named data streams.
<code>vxfs_nattr_link</code>	Links to a named data stream.
<code>vxfs_nattr_open</code>	Opens a named data stream.
<code>vxfs_nattr_rename</code>	Renames a named data stream.
<code>vxfs_nattr_unlink</code>	Removes a named data stream.
<code>vxfs_nattr_utimes</code>	Sets access and modification times for named data streams.
<code>vxfs_vol_add</code>	Adds a volume to a multi-volume file system.
<code>vxfs_vol_clearflags</code>	Clears specified flags on volumes in a multi-volume file system.
<code>vxfs_vol_deencapsulate</code>	De-encapsulates a volume from a multi-volume file system.
<code>vxfs_vol_encapsulate</code>	Encapsulates a volume within a multi-volume file system.
<code>vxfs_vol_encapsulate_bias</code>	Encapsulates a volume within a multi-volume file system.
<code>vxfs_vol_enumerate</code>	Returns information about the volumes within a multi-volume file system.
<code>vxfs_vol_queryflags</code>	Queries flags on volumes in a multi-volume file system.
<code>vxfs_vol_remove</code>	Removes a volume from a multi-volume file system.
<code>vxfs_vol_resize</code>	Resizes a specific volume within a multi-volume file system.
<code>vxfs_vol_setflags</code>	Sets specified flags on volumes in a multi-volume file system.
<code>vxfs_vol_stat</code>	Returns free space information about a component volume within a multi-volume file system.

[Table B-17](#) describes the VxFS-specific section 4 manual pages.

Table B-17 Section 4 manual pages

Section 4	Description
fs_vxfs	Provides the format of a VxFS file system volume.
inode_vxfs	Provides the format of a VxFS file system inode.
tunefstab	Describes the VxFS file system tuning parameters table.

[Table B-18](#) describes the VxFS-specific section 7 manual pages.

Table B-18 Section 7 manual pages

Section 7	Description
vxfsio	Describes the VxFS file system control functions.

Creating a starter database

This appendix includes the following topics:

- [Creating a database for Oracle 11gr2](#)

Creating a database for Oracle 11gr2

Create a database tablespace for Oracle 11gr2 on shared raw VxVM volumes

Before you begin, take note of the following prerequisites:

- CRS daemons must be running. To verify the status of CRS, enter:
For Oracle 11gr2,

```
# $CRS_HOME/bin/crsctl status resource -t
```
- Use the `ping` command to verify that all private IP addresses on each node are up.

Creating database tablespace on shared raw VxVM volumes

This section describes how to create database tablespace on shared raw VxVM volumes (option 1).

To create database tablespace on shared raw VxVM volumes (option 1)

- 1 On any cluster node, log in as `root`. Find out the spare disks that can be used for creating shared disk group for Oracle database tablespaces, enter:

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
sda	auto:none	-	-	online invalid
sdb	auto:none	-	-	online invalid
sdc	auto:cdsdisk	-	tempgd	online shared
sdd	auto:cfdisk	-	ocrvotgd	online shared
sde	auto:cdsdisk	-	-	online shared
sdf	auto:cdsdisk	-	-	online shared

The above sample output indicates that shared disks `sde` and `sdf` are free and can be used for Oracle database tablespaces.

- 2 Create a shared disk group:

```
# vxdg -s init oradatadg sde sdf
```

- 3 Create a volume in the shared group for each of the required tablespaces.

See the *Oracle* documentation specific to the Oracle database release to determine the tablespace requirements.

For example, enter:

```
# vxassist -g oradatadg make VRT_system01 1000M
# vxassist -g oradatadg make VRT_system02 10M
.
.
.
```

- 4 Define the access mode and permissions for the volumes storing the Oracle data. For each volume listed in `$ORACLE_HOME/raw_config`, use the `vxedit` command:

```
# vxedit -g disk_group set group=group user=user mode=660  
volume
```

See the `vxedit(1M)` manual page.

For example, enter:

```
# vxedit -g oradatadg set group=oinstall user=oracle mode=660 \  
VRT_system01
```

In this example, `VRT_system01` is the name of one of the volumes. Repeat the command to define access mode and permissions for each volume in the `oradatadg`.

- 5 Create the database.

See the *Oracle* documentation.

Index

Symbols

- /dev/vx/dmp directory 172
- /dev/vx/rdmp directory 172
- /etc/default/vxassist file 190, 782
- /etc/default/vxdg defaults file 151
- /etc/init.d/vxvm-recover file 786
- /etc/vfstab file 361, 854
- /etc/volboot file 104
- /etc/vx/darecs file 104
- /etc/vx/dmppolicy.info file 292
- /kernel/drv/vxio.conf file 921
- /kernel/drv/vxspec.conf file 921
- /lib/svc/method/vxvm-recover file 786

A

- A/A disk arrays 170
- A/P disk arrays 171
- A/P-C disk arrays 171–172
- A/P-F disk arrays 171
- A/P-G disk arrays 172
- A5x00 arrays
 - removing and replacing disks 864
- access control lists 82
- access port 171
- Actions 513
- activation modes for shared disk groups 150–151
- active path attribute 289
- active paths
 - devices 290
- ACTIVE state 589
- Active/Active disk arrays 170
- Active/Passive disk arrays 171
- adaptive load-balancing 292
- adding disks 345
- Administering
 - Clustered NFS 518
- Administration
 - I/O fencing 390
- administration
 - SFCFSHA 355

Agents

- CVM 74
- allocation
 - site-based 465
- allocation policies 236
 - default 236
 - extent 78
 - extent based 78
- APM
 - configuring 307
- Architecture
 - SFCFSHA 41
- array policy module (APM)
 - configuring 307
- array ports
 - disabling for DMP 298
 - displaying information about 278
 - enabling for DMP 299
- array support library (ASL) 246
- Array Volume ID
 - device naming 328
- arrays
 - DMP support 245
- ASL
 - array support library 245–246
- Asymmetric mounts 46
- asymmetric mounts
 - mount_vxfs(1M) 47
- ATTACHING state 589
- attributes
 - active 289
 - autogrow 598, 602
 - autogrowby 598
 - cache 602
 - cachesize 602
 - dcolon 140, 636
 - for specifying storage 204
 - highwatermark 598
 - maxautogrow 598
 - mirdg 610
 - mirvol 610
 - ncachemirror 602

attributes *(continued)*

- ndcomirror 636
- ndcomirs 597
- newvol 608
- nmirror 608
- nomanual 289
- nopreferred 289
- preferred priority 289
- primary 289
- secondary 289
- setting for paths 289, 291
- snapvol 604, 610
- source 604, 610
- standby 289
- syncing 595, 622

autogrow

- tuning 624
- autogrow attribute 598, 602
- autogrowby attribute 598
- autotrespass mode 171
- availability
 - storage 152

B

Backup strategies

- SFCFSHA 49

backups

- created using snapshots 595
- creating for volumes 567
- creating using instant snapshots 595
- creating using third-mirror snapshots 626
- for multiple volumes 611, 631

bad block revectoring 218

balanced path policy 293

Basic design

- Clustered NFS 511

blkclear 80

blkclear mount option 218

block based architecture 87

blockmap for a snapshot file system 669

blocks on disks 100

boot command

- a flag 878
- syntax 878

boot disk

- encapsulating 875
- listing volumes on 879
- mirroring 875
- unencapsulating 883

boot disks

- alternate 876

boot-time restrictions 873

booting root volumes 873

BROKEN state 589

buffered file systems 80

buffered I/O 487

C

c# 94

cache

- for space-optimized instant snapshots 569

cache advisories 489

cache attribute 602

cache objects

- creating 598
- enabling 599
- listing snapshots in 624

caches

- creating 598
- deleting 626
- finding out snapshots configured on 626
- growing 625
- listing snapshots in 624
- removing 626
- resizing 625
- shrinking 625
- stopping 626
- used by space-optimized instant snapshots 570

cachesize attribute 602

Campus Cluster feature

- administering 464

campus clusters

- administering 464

cascade instant snapshots 590

cascaded snapshot hierarchies

- creating 615

categories

- disks 246

CFS file system

- Sharing 518, 546

CFS primaryship

- determining 47
- moving 47

cfscluster command 359

cfsdgadm command 359

cfsmntadm command 359

cfsmount command 359

- cfsnfssg_dummy
 - service group 515
- cfsshare
 - manual page 513
- cfsshare (1M) 538
- cfsshare command 359
- cfsumount command 359
- Changing the CVM master 373
- check_all policy 305
- check_alternate policy 306
- check_disabled policy 306
- check_periodic policy 306
- checkpoint interval 893
- CIFS
 - Common Internet File System 537
 - debugging 553
- cio
 - Concurrent I/O 223
- closesync 80
- Cluster file systems
 - VxFS
 - unsupported features 45
- cluster file systems
 - support features
 - VxFS 42
- Cluster master node
 - changing 373
- cluster mount 84
- cluster snapshot
 - characteristics 362
- cluster-shareable disk groups in clusters 149
- Clustered NFS 510
 - Administering 518
 - Basic design 511
 - Configure 513
 - Debugging 535
 - Functionality 511
 - Unconfigure 515
 - Understanding 511
- clusters
 - activating disk groups 151
 - activating shared disk groups 384
 - activation modes for shared disk groups 150
 - benefits 145
 - checking cluster protocol version 386
 - cluster-shareable disk groups 149
 - configuration 158
 - configuring exclusive open of volume by node 385
 - clusters (*continued*)
 - connectivity policies 152
 - converting shared disk groups to private 383
 - creating shared disk groups 381
 - designating shareable disk groups 149
 - detach policies 152
 - determining if disks are shared 380
 - forcibly adding disks to disk groups 382
 - forcibly importing disk groups 382
 - importing disk groups as shared 382
 - initialization 158
 - introduced 146
 - limitations of shared disk groups 157
 - listing shared disk groups 380
 - maximum number of nodes in 146
 - moving objects between disk groups 383
 - node shutdown 164
 - nodes 147
 - operation of DRL in 165–166
 - operation of vxconfigd in 162
 - operation of VxVM in 147
 - private disk groups 149
 - private networks 148
 - protection against simultaneous writes 150
 - reconfiguration of 159
 - resolving disk status in 152
 - shared disk groups 149
 - shared objects 150
 - splitting disk groups in 383
 - use of DMP in 180
 - vol_fmr_logsz tunable 899
 - volume reconfiguration 161
 - vxclustadm 159
 - vxctl 379
 - vxrecover 387
 - vxstat 387
 - CNFS 510
 - columns
 - changing number of 835
 - in striping 109
 - mirroring in striped-mirror volumes 202
- CommandCentral Storage 142
- commands
 - cfscluster 359
 - cfsdgadm 359
 - cfsmntadm 359
 - cfsmount 359
 - cfsshare 359
 - cfsumount 359

- commands (*continued*)
 - cron 90
 - fsadm 90
 - getext 239
 - setext 239
 - Common Internet File System
 - CIFS 537
 - compressing files 87
 - concatenated volumes 106, 198
 - concatenated-mirror volumes
 - creating 201
 - defined 114
 - recovery 199
 - concatenation 106
 - configuration changes
 - monitoring using vxnotify 831
 - configuration database
 - reducing size of 814
 - Configure
 - Clustered NFS 513
 - Configuring
 - low priority link 52
 - Configuring a CNFS
 - samples 522, 525
 - Configuring DMP
 - using templates 907
 - connectivity policies 152
 - contiguous reservation 238
 - Controller ID
 - displaying 277
 - controllers
 - disabling for DMP 298
 - disabling in DMP 265
 - displaying information about 276
 - enabling for DMP 299
 - mirroring across 208
 - number 94
 - specifying to vxassist 204
 - converting a data Storage Checkpoint to a nodata Storage Checkpoint 645
 - convosync mount option 215, 219
 - coordinator disks
 - DMP devices 57
 - for I/O fencing 57
 - copy-on-write
 - used by instant snapshots 588
 - copy-on-write technique 574, 640
 - copymaps 138–139
 - CP server
 - deployment scenarios 430
 - migration scenarios 430
 - CP server database 69
 - CP server user privileges 69
 - cp_vxfs 239
 - cpio_vxfs 239
 - Creating
 - database for Oracle 11gr2 986
 - creating
 - snapshot
 - SFCFSHA 363
 - creating a multi-volume support file system 712
 - creating file systems with large files 222
 - creating files with mkfs 211, 214
 - cron 90, 231
 - cron sample script 233
 - customized naming
 - DMP nodes 329
 - CVM
 - agents 74
 - cluster functionality of VxVM 145
 - failure tolerance 152
 - functionality 74
 - resilience 152
 - storage disconnectivity tolerance 152
 - CVM master
 - changing 373
- ## D
- d# 94
 - data change object
 - DCO 138
 - data copy 486
 - data corruption
 - preventing 55
 - data deduplication 86
 - data disks
 - for I/O fencing 57
 - data integrity 80
 - data redundancy 112, 115
 - data Storage Checkpoints definition 578
 - data synchronous I/O 219, 487
 - data transfer 486
 - data volume configuration 129
 - database replay logs and sequential DRL 128
 - databases
 - integrity of data in 568
 - resilvering 128

- databases (*continued*)
 - resynchronizing 128
- DCO
 - adding version 0 DCOs to volumes 635
 - considerations for disk layout 820
 - data change object 138
 - dissociating version 0 DCOs from volumes 638
 - effect on disk group split and join 820
 - instant snap version 138
 - log plexes 136
 - log volume 138
 - moving log plexes 637
 - reattaching version 0 DCOs to volumes 638
 - removing version 0 DCOs from volumes 638
 - specifying storage for version 0 plexes 637
 - used with DRL 128
 - version 0 138
 - version 20 138–139
 - versioning 138
- dcolon attribute 140, 636
- DDI_NT_FABRIC property 244
- DDL 174
 - Device Discovery Layer 249
- debugging CIFS 553
- Debugging Clustered NFS 535
- deduplicating data 86
- default
 - allocation policy 236
- defaultdg 192
- defragmentation 90
 - extent 232
 - scheduling with cron 232
- delaylog mount option 216
- detach policy
 - global 154
 - local 155
- Determining
 - CFS primaryship 47
- device discovery
 - introduced 174
 - partial 244
- Device Discovery Layer 249
- Device Discovery Layer (DDL) 174, 249
- device names 94
 - configuring persistent 330
 - user-specified 329
- device nodes
 - controlling access for volume sets 707
 - displaying access for volume sets 707
- device nodes (*continued*)
 - enabling access for volume sets 706
 - for volume sets 705
- devices
 - adding foreign 261
 - fabric 244
 - JBOD 245
 - listing all 251
 - making invisible to VxVM 262
 - nopriv 871
 - path redundancy 290
- dgfailpolicy attribute 384
- direct data transfer 486
- direct I/O 486
- directory reorganization 233
- dirty flags set on volumes 126
- dirty region logging.. See DRL
- dirty regions 900
- disable failure policy 154
- disabled file system
 - snapshot 584
- disabled paths 268
- discovered direct I/O 487
- disk access records
 - stored in /etc/vx/darecs 104
- disk arrays
 - A/A 170
 - A/P 171
 - A/P-F 171
 - A/P-G 172
 - Active/Active 170
 - Active/Passive 171
 - adding disks to DISKS category 258
 - defined 95
 - excluding support for 257
 - JBOD devices 245
 - listing excluded 257
 - listing supported 256
 - listing supported disks in DISKS category 257
 - multipathed 104
 - re-including support for 257
 - removing disks from DISKS category 260
 - supported with DMP 256
- Disk Group Split/Join 570
- disk groups
 - activating shared 384
 - activation in clusters 151
 - cluster-shareable 149
 - configuring site consistency on 472

disk groups *(continued)*

- configuring site-based allocation on 471
- converting to private 383
- creating shared 381
- creating with old version number 849
- defaults file for shared 151
- defined 98
- designating as shareable 149
- determining the default disk group 191
- displaying boot disk group 191
- displaying default disk group 192
- displaying version of 849
- failure policy 154
- features supported by version 844
- free space in 777
- importing as shared 382
- importing forcibly 382
- joining 816, 825
- layout of DCO plexes 820
- limitations of move
 - split. *See* and join
- listing objects affected by a move 819
- listing shared 380
- making site consistent 470
- moving disks between 813, 822
- moving licensed EMC disks between 822
- moving objects between 814, 821
- moving objects in clusters 383
- private in clusters 149
- recovery from failed reconfiguration 818
- reorganizing 814
- restarting moved volumes 826
- root 99
- rootdg 99
- setting default disk group 192
- shared in clusters 149
- splitting 815, 824
- splitting in clusters 383
- upgrading version of 848–849
- version 844, 848

disk layout

- Version 1 939
- Version 2 939
- Version 3 939
- Version 4 940
- Version 5 940
- Version 6 940
- Version 7 940
- Version 8 940

disk layout *(continued)*

- Version 9 940
- Disk layout versions 43
- disk media names 99
- disk names
 - configuring persistent 330
- disk### 100
- disk###-### 100
- diskdetpolicy attribute 384
- disks 246
 - adding 345
 - adding to disk groups forcibly 382
 - adding to DISKS category 258
 - array support library 246
 - categories 246
 - changing naming scheme 327
 - complete failure messages 776
 - configuring newly added 243
 - configuring persistent names 330
 - determining failed 776
 - determining if shared 380
 - Device Discovery Layer 249
 - disabled path 268
 - discovery of by DMP 243
 - discovery of by VxVM 245
 - disk access records file 104
 - disk arrays 95
 - displaying information 325–326
 - displaying information about 325
 - displaying naming scheme 328
 - displaying spare 779
 - dynamic LUN expansion 318
 - enabled path 268
 - encapsulation 866, 872
 - enclosures 174
 - excluding free space from hot-relocation use 781
 - failure handled by hot-relocation 773
 - formatting 336
 - hot-relocation 771
 - initializing 336
 - installing 336
 - invoking discovery of 247
 - layout of DCO plexes 820
 - listing those supported in JBODs 257
 - making available for hot-relocation 779
 - making free space available for hot-relocation use 782
 - marking as spare 779
 - mirroring boot disk 875

disks *(continued)*

- mirroring root disk 875
- mirroring volumes on 839
- moving between disk groups 813, 822
- moving volumes from 812
- nopriv devices 871
- number 94
- OTHER_DISKS category 246
- partial failure messages 775
- postponing replacement 858
- primary path 268
- reinitializing 344
- removing 345, 858
- removing from A5x00 arrays 864
- removing from DISKS category 260
- removing from pool of hot-relocation spares 780
- removing from VxVM control 855
- removing with subdisks 347
- renaming 348
- replacing 858
- replacing in A5x00 arrays 864
- replacing removed 862
- resolving status in clusters 152
- root disk 872
- scanning for 244
- secondary path 268
- spare 777
- specifying to vxassist 204
- tagging with site name 475
- VM 99

DISKS category 246

- adding disks 258
- listing supported disks 257
- removing disks 260

displaying

- DMP nodes 272
- HBA information 277
- redundancy levels 290
- supported disk arrays 256

displaying mounted file systems 229**displaying statistics**

- erroneous I/Os 286
- queued I/Os 286

DMP

- booting from DMP devices 180
- check_all restore policy 305
- check_alternate restore policy 306
- check_disabled restore policy 306
- check_periodic restore policy 306

DMP *(continued)*

- configuring disk devices 243
- configuring DMP path restoration policies 305
- configuring I/O throttling 302
- configuring response to I/O errors 301, 304
- disabling array ports 298
- disabling controllers 298
- disabling multi-pathing 262
- disabling paths 298
- disk discovery 243
- displaying DMP database information 265
- displaying DMP node for a path 271
- displaying DMP node for an enclosure 271–272
- displaying DMP nodes 272
- displaying information about array ports 278
- displaying information about controllers 276
- displaying information about enclosures 277
- displaying information about paths 266
- displaying LUN group for a node 273
- displaying paths controlled by DMP node 274
- displaying paths for a controller 275
- displaying paths for an array port 275
- displaying recoveryoption values 304
- displaying status of DMP path restoration
 - thread 307
- displaying TPD information 279
- dynamic multi-pathing 170
- enabling array ports 299
- enabling controllers 299
- enabling multi-pathing 264
- enabling paths 299
- enclosure-based naming 173
- gathering I/O statistics 283
- in a clustered environment 180
- load balancing 178
- logging levels 886
- metanodes 172
- nodes 172
- path aging 886
- path failover mechanism 177
- path-switch tunable 889
- renaming an enclosure 300
- restore policy 305
- scheduling I/O on secondary paths 295
- setting the DMP restore polling interval 305
- stopping the DMP restore daemon 307
- tuning with templates 907
- vxdmpadm 269

- DMP nodes
 - displaying consolidated information 272
 - setting names 329
- DMP support
 - JBOD devices 245
- dmp_cache_open tunable 885
- dmp_daemon_count tunable 885
- dmp_delayq_interval tunable 885
- dmp_fast_recovery tunable 885
- dmp_health_time tunable 886
- dmp_log_level tunable 886
- dmp_low_impact_probe 886
- dmp_lun_retry_timeout tunable 887
- dmp_monitor_fabric tunable 887
- dmp_monitor_osevent tunable 888
- dmp_monitor_ownership tunable 888
- dmp_native_multipathing tunable 888
- dmp_native_support tunable 889
- dmp_path_age tunable 889
- dmp_pathswitch_blks_shift tunable 889
- dmp_probe_idle_lun tunable 890
- dmp_probe_threshold tunable 890
- dmp_restore_cycles tunable 890
- dmp_restore_interval tunable 890
- dmp_restore_state tunable 891
- dmp_scsi_timeout tunable 891
- dmp_sfg_threshold tunable 892
- dmp_stat_interval tunable 892
- DR
 - dynamic reconfiguration 179
- DRL
 - dirty region logging 127
 - handling recovery in clusters 166
 - hot-relocation limitations 773
 - log subdisks 128
 - maximum number of dirty regions 900
 - minimum number of sectors 900
 - operation in clusters 165
 - sequential 128
 - use of DCO with 128
- dynamic LUN expansion 318
- dynamic reconfiguration 179

E

- EMC arrays
 - moving disks between disk groups 822
- EMC PowerPath
 - coexistence with DMP 248

- EMC Symmetrix
 - autodiscovery 248
- enabled paths
 - displaying 268
- encapsulating disks 866, 872
- encapsulating volumes 710
- encapsulation
 - failure of 870
- enclosure-based naming 174, 327
 - DMP 173
- enclosures 174
 - displaying information about 277
 - issues with nopriv disks 332
 - issues with simple disks 332
 - path redundancy 290
 - setting attributes of paths 289, 291
 - tagging with site name 476, 479
- enhanced data integrity modes 80
- ENOSPC 654
- Environment
 - public network 52
 - single private link 52
- erroneous I/Os
 - displaying statistics 286
- error messages
 - Association count is incorrect 168
 - Association not resolved 168
 - Cannot auto-import group 168
 - Configuration records are inconsistent 168
 - Disk group has no valid configuration copies 169
 - Disk group version doesn't support feature 844
 - Disk not moving
 - but subdisks on it are 819
 - Duplicate record in configuration 168
 - It is not possible to encapsulate 870
 - The encapsulation operation failed 870
 - tmpsize too small to perform this relayout 122
 - unsupported layout 870
 - vxdg listmove failed 819
- error daemon 176
- errors
 - handling transient errors 887
- exclusive-write mode 151
- exclusivewrite mode 150
- expansion 90
- explicit failover mode 171
- extent 78, 235
 - attributes 235
 - reorganization 233

extent allocation 77–78

aligned 236

control 235

fixed size 235

external quotas file 924

F

fabric devices 244

FAILFAST flag 177

Failover

primary 42

secondary 42

failover 145–146

failover mode 171

failure handled by hot-relocation 773

failure in RAID-5 handled by hot-relocation 773

failure policies 154

failures

storage 152

Fast file system recovery 43

FastResync

effect of growing volume on 140

limitations 141

Non-Persistent 134

Persistent 134–135, 567

size of bitmap 899

snapshot enhancements 586

use with snapshots 133

FastResync/cache object metadata cache size

tunable 903

fc_off 933

file

sparse 237

file compression 87

file system

block size 241

buffering 80

displaying mounted 229

mount 532

file system tunables

tunefstab(4) 48

vxtunefs(1M) 48

file systems

mirroring on root disk 878

unmounting 854

filesystem

primary 572

FileSnap 44

FileSnaps

about 580

data mining, reporting, and testing 664

virtual desktops 663

write intensive applications 664

backup 582

best practices 663

block map fragmentation 582

concurrent I/O 581

copy-on-write 581

creation 661

properties 580

reading from 582

using 662

fire drill

defined 465

testing 472

fire drills

about 459

disaster recovery 459

for global clusters 459

fixed extent size 235

fixed write size 237

FlashSnap 565

FMR.. See FastResync

foreign devices

adding 261

formatting disks 336

fragmentation

monitoring 231, 233

reorganization facilities 231

reporting 231

fragmented file system characteristics 232

free space in disk groups 777

free space monitoring 231

Freeze 44

freeze 489

freezing and thawing, relation to Storage

Checkpoints 572

fsadm 90

how to minimize file system free space

fragmentation 228

how to reorganize a file system 228

how to resize a file system 226

reporting extent fragmentation 232

scheduling defragmentation using cron 232

fsadm_vxfs 223

fsadm_vxfs(1M)

manual page 360

- fsckat 668
- fsck 645
- fsckptadm
 - Storage Checkpoint administration 641
- fsclustadm 361, 511
- fsclustadm(1M)
 - manual page 360
- fstyp
 - how to determine the file system type 230
- fsvoladm 712
- full-sized instant snapshots 587
 - creating 604
 - creating volumes for use as 600
- fullinst snapshot type 621
- Functionality
 - Clustered NFS 511

G

- GAB 158
- get I/O parameter ioctl 490
- getext 239
- getfacl 82
- GLM 42
 - SFCFSHA 45
- global detach policy 154
- Group Membership and Atomic Broadcast (GAB) 158
- growing
 - file system
 - SFCFSHA 361

H

- handling storage disconnectivity
 - CVM 152
- HBA information
 - displaying 277
- HBAs
 - listing ports 252
 - listing supported 251
 - listing targets 252
- highwatermark attribute 598
- host failures 481
- hot-relocation
 - complete failure messages 776
 - configuration summary 778
 - daemon 772
 - defined 127
 - detecting disk failure 773
 - detecting plex failure 773

- hot-relocation (*continued*)
 - detecting RAID-5 subdisk failure 773
 - excluding free space on disks from use by 781
 - limitations 773
 - making free space on disks available for use
 - by 782
 - marking disks as spare 779
 - modifying behavior of 786
 - notifying users other than root 787
 - operation of 771
 - partial failure messages 775
 - preventing from running 787
 - reducing performance impact of recovery 787
 - removing disks from spare pool 780
 - subdisk relocation 778
 - subdisk relocation messages 783
 - unrelocating subdisks 783
 - unrelocating subdisks using vxunreloc 784
 - use of free space in disk groups 777
 - use of spare disks 777
 - use of spare disks and free space 777
 - using only spare disks for 782
 - vxrelocd 772
- how to access a Storage Checkpoint 643
- how to create a Storage Checkpoint 642
- how to determine the file system type 230
- how to display mounted file systems 225
- how to minimize file system free space
 - fragmentation 228
- how to mount a Storage Checkpoint 643
- how to remove a Storage Checkpoint 643
- how to reorganize a file system 228
- how to resize a file system 226
- how to unmount a Storage Checkpoint 645

I

- I/O
 - direct 486
 - gathering statistics for DMP 283
 - kernel threads 93
 - scheduling on secondary paths 295
 - sequential 487
 - synchronous 487
 - throttling 177
- I/O error handling 51
- I/O failures
 - Recovering 51
- I/O fencing
 - administration 390

- I/O fencing (*continued*)
 - operations 56
 - preventing data corruption 55
 - testing and scenarios 62
- I/O operations
 - maximum size of 894
- I/O policy
 - displaying 291
 - example 296
 - specifying 292
- I/O requests
 - asynchronous 219
 - synchronous 218
- I/O throttling 302
- I/O throttling options
 - configuring 305
- identifiers for tasks 828
- idle LUNs 890
- implicit failover mode 171
- initialization
 - of disks 336
- inodes, block based 78
- instant snap version
 - of DCOs 138
- instant snapshots
 - backing up multiple volumes 611
 - cascaded 590
 - creating backups 595
 - creating for volume sets 612
 - creating full-sized 604
 - creating space-optimized 601
 - creating volumes for use as full-sized 600
 - displaying information about 620
 - dissociating 618
 - full-sized 587
 - improving performance of synchronization 623
 - reattaching 616
 - refreshing 616
 - removing 619
 - restoring volumes using 618
 - space-optimized 569
 - splitting hierarchies 619
 - synchronizing 622
- intent log 78
 - multi-volume file system support 710
- intent log resizing 79
- intent logging 568
- internal quotas file 924
- ioctl calls 894–895

- ioctl interface 235
- iSCSI parameters
 - administering with DDL 254
 - setting with vxddladm 254

J

- JBOD
 - DMP support 245
- JBODs
 - adding disks to DISKS category 258
 - listing supported disks 257
 - removing disks from DISKS category 260
- Jeopardy 53
 - handling 53
 - state 53

L

- large files 81, 222
 - creating file systems with 222
 - mounting file systems with 223
- largefiles mount option 223
- layered volumes
 - defined 119, 199
 - striped-mirror 113
- layouts
 - left-symmetric 117
 - types of volume 198
- leave failure policy 154
- left-symmetric layout 117
- link objects 589
- linked break-off snapshots 589
 - creating 609
- linked third-mirror snapshots
 - reattaching 617
- listing
 - alternate boot disks 876
 - DMP nodes 272
 - supported disk arrays 256
- load balancing 170
 - across nodes in a cluster 146
 - displaying policy for 291
 - specifying policy for 292
- local detach policy 155
- local mount 84
- Locking 44
- log mount option 215–216
- log subdisks
 - DRL 128

- logdisk 203–204
- logical units 171
- logiosize mount option 217
- logs
 - RAID-5 119, 126
 - specifying number for RAID-5 203
- Low priority link
 - configuring 52
- LUN 171
- LUN expansion 318
- LUN group failover 172
- LUN groups
 - displaying details of 273
- LUNs
 - idle 890
 - thin provisioning 679
- M**
- main.cf file 526
 - sample 547
- major and minor numbers
 - checking 516
 - shared devices 515
- Manual page
 - fsadm_vxfs(1M) 360
 - fsclustadm(1M) 360
 - mount(1M) 359
- manual page
 - cfsshare 513
- Master node
 - changing 373
- master node
 - defined 148
 - discovering 379
- maxautogrow attribute 598
- memory
 - granularity of allocation by VxVM 896
 - maximum size of pool for VxVM 896
 - minimum size of pool for VxVM 898
 - persistence of FastResync in 134
- Memory mapping 44
- messages
 - complete disk failure 776
 - hot-relocation of subdisks 783
 - partial disk failure 775
- metadata
 - multi-volume support 710
- metanodes
 - DMP 172
- migrating to thin storage 679
- mincache mount option 215, 218
- minimum queue load balancing policy 294
- minimum redundancy levels
 - displaying for a device 290
 - specifying for a device 290
- mirbrk snapshot type 621
- mirdg attribute 610
- mirrored volumes
 - changing read policies for 209
 - configuring VxVM to create by default 838
 - creating 200
 - creating across controllers 208
 - creating across targets 206
 - defined 198
 - dirty region logging 127
 - DRL 127
 - FastResync 127
 - FR 127
 - logging 127
 - snapshots 133
- mirrored-concatenated volumes
 - creating 200
 - defined 112
- mirrored-stripe volumes
 - benefits of 112
 - creating 202
 - defined 199
- mirroring
 - boot disk 875
 - defined 112
 - root disk 875
- mirroring plus striping 113
- mirrors
 - adding to volumes 838
 - creating snapshot 628
 - defined 103
 - removing from volumes 842
 - specifying number of 200
- mirvol attribute 610
- mirvol snapshot type 621
- mkfs
 - creating files with 211, 214
 - creating large files 223
- modes
 - enhanced data integrity 80
- monitoring fragmentation 231
- mount 79, 223
 - how to display mounted file systems 225

- mount (*continued*)
 - mounting a Storage Checkpoint 643
 - nfs-exported file system 532
 - pseudo device 644
- mount options 213
 - blkclear 218
 - choosing 213
 - combining 224
 - convosync 215, 219
 - delaylog 81, 216
 - extended 79
 - largefiles 223
 - log 80, 215–216
 - logiosize 217
 - mincache 215, 218
 - nodatainlog 215, 218
 - tmplog 216
- mount(1M)
 - manual page 359
- mount_vxfs(1M)
 - asymmetric mounts 47
- mounted file system
 - displaying 229
- mounting a file system
 - option combinations 224
 - with large files 223
- mounting a Storage Checkpoint 645
- mounting a Storage Checkpoint of a cluster file system 645
- Moving
 - CFS primaryship 47
- multi-pathing
 - disabling 262
 - displaying information about 266
 - enabling 264
- multi-volume file systems 709
- Multi-Volume Support 701
- multi-volume support
 - creating a MVS file system 712
- multiple block operations 78
- mv_vxfs 239

N

- name space
 - preserved by Storage Checkpoints 641
- names
 - defining for snapshot volumes 631
 - device 94
 - disk media 99

- names (*continued*)
 - plex 102
 - renaming disks 348
 - subdisk 100
 - VM disk 100
 - volume 102
- naming
 - DMP nodes 329
- naming scheme
 - changing for disks 327
 - changing for TPD enclosures 331
 - displaying for disks 328
- ncachemirror attribute 602
- ncheck 937
- ndcomirror attribute 636
- ndcomirs attribute 597
- Nested mounts 44
- newvol attribute 608
- NFS mounts 44
- NFS services
 - shared storage 515
- nfs-exported file system
 - mount 532
- nmirror attribute 607–608
- nodata Storage Checkpoints 645
- nodata Storage Checkpoints definition 578
- nodatainlog mount option 215, 218
- nodes
 - DMP 172
 - in clusters 147
 - maximum number in a cluster 146
 - requesting status of 379
 - shutdown in clusters 164
 - use of vxclustadm to control CVM functionality 159
- nommanual path attribute 289
- non-autotrespass mode 171
- Non-Persistent FastResync 134
- nopreferred path attribute 289
- nopriv devices 871
- nopriv disks
 - issues with enclosures 332
- NTP
 - network time protocol daemon 48, 360

O

- O_SYNC 215
- objects
 - physical 94

- objects (*continued*)
 - virtual 96
- off-host processing 146
- OMF 498
 - working with Oracle Disk Manager 498
- online invalid status 326
- online relayout
 - changing number of columns 835
 - changing region size 837
 - changing speed of 837
 - changing stripe unit size 835
 - controlling progress of 837
 - defined 121
 - destination layouts 831
 - failure recovery 125
 - how it works 121
 - limitations 124
 - monitoring tasks for 836
 - pausing 837
 - performing 831
 - resuming 837
 - reversing direction of 837
 - specifying non-default 835
 - specifying plexes 836
 - specifying task tags for 836
 - temporary area 122
 - transformation characteristics 125
 - transformations and volume length 125
 - types of transformation 832
 - viewing status of 836
- online status 326
- OpenBoot PROMs (OPB) 878
- Oracle Disk Manager 494
 - benefits 495
 - converting Quick I/O files 501
 - disabling 505
 - migrating files to 501
 - preparing existing databases for use with 501
 - setting up 500
- Oracle Managed Files 498
 - working with Oracle Disk Manager 498
- ordered allocation 204, 206
- OTHER_DISKS category 246

P

- Parallel fsck threads 48
- Parallel I/O 50
- parity in RAID-5 115
- partial device discovery 244
- partition size
 - displaying the value of 291
 - specifying 293
- partitioned directories 86
- partitions
 - number 94–95
 - slices 95
- path aging 886
- path failover in DMP 177
- pathgroups
 - creating 263
- paths
 - disabling for DMP 298
 - enabling for DMP 299
 - setting attributes of 289, 291
- performance
 - changing values of tunables 906, 914
 - displaying tunable values 922
 - improving for instant snapshot
 - synchronization 623
 - load balancing in DMP 178
 - overall 213
 - SFCFSHA 363
 - snapshot file systems 668
- persistence
 - device naming option 328
- persistent device name database 330
- persistent device naming 330
- Persistent FastResync 134–135, 140, 567
- physical disks
 - complete failure messages 776
 - determining failed 776
 - displaying information 325
 - displaying information about 325
 - displaying spare 779
 - excluding free space from hot-relocation use 781
 - failure handled by hot-relocation 773
 - installing 336
 - making available for hot-relocation 779
 - making free space available for hot-relocation
 - use 782
 - marking as spare 779
 - moving between disk groups 813, 822
 - moving volumes from 812
 - partial failure messages 775
 - postponing replacement 858
 - removing 345, 858
 - removing from A5x00 arrays 864
 - removing from pool of hot-relocation spares 780

- physical disks *(continued)*
 - removing with subdisks 347
 - replacing 858
 - replacing in A5x00 arrays 864
 - replacing removed 862
 - spare 777
- physical objects 94
- ping-pong effect 180
- plex attribute 608
- plexes
 - adding to snapshots 633
 - changing read policies for 209
 - complete failure messages 776
 - converting to snapshot 630
 - defined 101
 - failure in hot-relocation 773
 - maximum number of subdisks 896
 - maximum number per volume 102
 - mirrors 103
 - moving 637
 - names 102
 - partial failure messages 775
 - putting online 852
 - reattaching 852
 - recovering after correctable hardware failure 776
 - removing from volumes 842
 - sparse 125
 - specifying for online relayout 836
 - striped 108
 - types 101
- point-in-time copy solutions
 - applications 556
- polling interval for DMP restore 305
- ports
 - listing 252
- postoffline
 - trigger 512
- postonline
 - trigger 512
- PowerPath
 - coexistence with DMP 248
- prefer read policy 210
- preferred plex
 - read policy 210
- preferred priority path attribute 289
- preonline
 - trigger 511
- Primary
 - failover 42

- primary and secondary
 - terminology 47
- primary boot disk failure 876
- primary fails
 - SFCFSHA 361
- primary fileset relation to Storage Checkpoints 572
- primary path 171, 268
- primary path attribute 289
- primaryship
 - setting with fsclustadm 362
- priority load balancing 294
- private disk groups
 - converting from shared 383
 - in clusters 149
- private network
 - in clusters 148
- PROMs
 - boot 878
- prtconf
 - displaying tunables 922
- pseudo device 644

Q

- queued I/Os
 - displaying statistics 286
- Quick I/O
 - converting files to Oracle Disk Manager 501
- Quick I/O for databases 44
- quota commands 925
- quotacheck 926
- Quotas 44
- quotas 923
 - exceeding the soft limit 924
 - hard limit 660, 923
 - soft limit 923
- quotas file 924
- quotas.grp file 924

R

- RAID-0 108
- RAID-0+1 112
- RAID-1 112
- RAID-1+0 113
- RAID-5
 - hot-relocation limitations 773
 - logs 119, 126
 - parity 115
 - specifying number of logs 203

- RAID-5 (*continued*)
 - subdisk failure handled by hot-relocation 773
 - volumes 115
- RAID-5 volumes
 - changing number of columns 835
 - changing stripe unit size 835
 - creating 203
 - defined 199
- raw device nodes
 - controlling access for volume sets 707
 - displaying access for volume sets 707
 - enabling access for volume sets 706
 - for volume sets 705
- read policies
 - changing 209
 - prefer 210
 - round 209
 - select 210
 - siteread 210, 466–467, 469
 - split 210
- read-only mode 151
- read-only Storage Checkpoints 643
- readonly mode 150
- Recovering
 - I/O failures 51
- recovery
 - checkpoint interval 893
 - I/O delay 893
- recovery accelerator 128
- recovery option values
 - configuring 304
- redo log configuration 129
- redundancy
 - of data on mirrors 198
 - of data on RAID-5 199
- redundancy levels
 - displaying for a device 290
 - specifying for a device 290
- redundant-loop access 176
- regionsize attribute 597–598
- reinitialization of disks 344
- relayout
 - changing number of columns 835
 - changing region size 837
 - changing speed of 837
 - changing stripe unit size 835
 - controlling progress of 837
 - limitations 124
 - monitoring tasks for 836
- relayout (*continued*)
 - online 121
 - pausing 837
 - performing online 831
 - resuming 837
 - reversing direction of 837
 - specifying non-default 835
 - specifying plexes 836
 - specifying task tags for 836
 - storage 121
 - transformation characteristics 125
 - types of transformation 832
 - viewing status of 836
- relocation
 - automatic 771
 - complete failure messages 776
 - limitations 773
 - partial failure messages 775
- Remote Mirror feature
 - administering 464
- remote mirrors
 - administering 464
- removable Storage Checkpoints definition 579
- removing devices
 - from VxVM control 262
- removing disks 858
- removing physical disks 345
- reorganization
 - directory 233
 - extent 233
- replacing disks 858
- replay logs and sequential DRL 128
- report extent fragmentation 231
- reservation space 235
- reservations
 - description 56
- resilvering
 - databases 128
- restore policy
 - check_all 305
 - check_alternate 306
 - check_disabled 306
 - check_periodic 306
- restored daemon 176
- resyncfromoriginal snapback 594
- resyncfromreplica snapback 594
- resynchronization
 - checkpoint interval 893
 - I/O delay 893

- resynchronization *(continued)*
 - of volumes 126
- resynchronizing
 - databases 128
 - snapshots 570
- retry option values
 - configuring 304
- Reverse Path Name Lookup 937

- root disk
 - defined 872
 - encapsulating 875
 - listing volumes on 879
 - mirroring 875
 - mirroring other file systems on 878
 - unencapsulating 883
- root disk group 99
- root disks
 - booting alternate 876
- root volume 873
- rootability 872
 - removing 883
- rootdg 99
- round read policy 209
- round-robin
 - load balancing 294
 - read policy 209

S

- s# 94–95
- sample
 - main.cf file 547
- SAN storage
 - configuring volumes on 142
- scandisks
 - vxdisk subcommand 244
- SCSI-3 PR 56
- Secondary
 - failover 42
- secondary path 171
- secondary path attribute 289
- secondary path display 268
- secure communication 70
- security 70
- select read policy 210
- sequential DRL
 - defined 128
 - maximum number of dirty regions 900
- sequential I/O 487
- serial split brain condition 465

- server-based fencing
 - replacing coordination points
 - online cluster 426
- service group
 - cfsnfssg_dummy 515
- setext 239
- setfacl 82
- Setting
 - parallel fsck threads 48
 - primaryship
 - fsclustadm 362
- setting
 - path redundancy levels 290
- SFCFSHA
 - administration 355
 - architecture 41
 - backup strategies 49
 - environments 52
 - features 40
 - GLM 45
 - growing
 - file system 361
 - load distribution 362
 - performance 363
 - primary fails 361
 - snapshots 49, 362
 - synchronize time 48
 - usage 40
- Shared CFS file system
 - unsharing 519
- shared CFS file system
 - unsharing 547
- shared disk groups
 - activating 384
 - activation modes 150–151
 - converting to private 383
 - creating 381
 - importing 382
 - in clusters 149
 - limitations of 157
 - listing 380
- shared storage
 - NFS services 515
- shared-read mode 151
- shared-write mode 151
- sharedread mode 150
- sharedwrite mode 150
- Sharing
 - CFS file system 518

- sharing
 - CFS file system 546
- simple disks
 - issues with enclosures 332
- single active path policy 295
- Site Awareness license 471
- site consistency
 - configuring 472
 - defined 465
- site failure
 - simulating 473
- site failures
 - host failures 481
 - loss of connectivity 480
 - recovery from 473, 481
 - scenarios and recovery procedures 479
 - storage failures 481
- site-based allocation
 - configuring for disk groups 471
 - defined 465
- site-based consistency
 - configuring on existing disk groups 470
- siteconsistent attribute 472
- siteread read policy 210, 466–467, 469
- sites
 - reattaching 473
- slave nodes
 - defined 148
- slices
 - number 94
 - partitions 95
- SmartMove feature 679
- SmartSync 128
 - disabling on shared disk groups 904
 - enabling on shared disk groups 904
- SmartTier 701
 - multi-volume file system support 710
- snap objects 137
- snap volume naming 594
- snapabort 586
- snapback
 - defined 587
 - merging snapshot volumes 632
 - resyncfromoriginal 594
 - resyncfromreplica 594, 632
- snapclear
 - creating independent volumes 633
- snapmir snapshot type 621
- snapof 671
- snapped file systems 82, 583
 - performance 668
 - unmounting 583
- snapread 668
- snapshot
 - characteristics
 - cluster 362
 - creating
 - SFCFSHA 363
- snapshot file systems 82, 583
 - blockmap 669
 - creating 671
 - data block area 669
 - disabled 584
 - fscat 668
 - fuser 583
 - mounting 671
 - multiple 583
 - on cluster file systems 583
 - performance 668
 - read 668
 - super-block 669
- snapshot hierarchies
 - creating 615
 - splitting 619
- snapshot mirrors
 - adding to volumes 614
 - removing from volumes 615
- Snapshots 44
 - SFCFSHA 49, 362
- snapshots
 - adding mirrors to volumes 614
 - adding plexes to 633
 - and FastResync 133
 - backing up multiple volumes 611, 631
 - backing up volumes online using 595
 - cascaded 590
 - comparison of features 131
 - converting plexes to 630
 - creating a hierarchy of 615
 - creating backups using third-mirror 626
 - creating for volume sets 612
 - creating full-sized instant 604
 - creating independent volumes 633
 - creating instant 595
 - creating linked break-off 609
 - creating snapshots of 591
 - creating space-optimized instant 601
 - creating third-mirror break-off 606

snapshots *(continued)*

- creating volumes for use as full-sized instant 600
- defining names for 631
- displaying information about 634
- displaying information about instant 620
- dissociating instant 618
- finding out those configured on a cache 626
- full-sized instant 132, 587
- hierarchy of 590
- improving performance of synchronization 623
- instant 569
- linked break-off 589
- listing for a cache 624
- merging with original volumes 632
- of volumes 130
- on multiple volumes 594
- reattaching instant 616
- reattaching linked third-mirror 617
- refreshing instant 616
- removing 630
- removing instant 619
- removing linked snapshots from volumes 615
- removing mirrors from volumes 615
- restoring from instant 618
- resynchronization on snapback 594
- resynchronizing 570
- resynchronizing volumes from 632
- space-optimized instant 569
- synchronizing instant 622
- third-mirror 131
- use of copy-on-write mechanism 588

snapsize 670
 snapstart 586
 snapvol attribute 604, 610
 snapwait 607, 610
 source attribute 604, 610
 space-optimized instant snapshots 569

- creating 601

 spaceopt snapshot type 621
 spanned volumes 106
 spanning 106
 spare disks

- displaying 779
- marking disks as 779
- used for hot-relocation 777

 sparse file 237
 sparse plexes 125
 specifying

- redundancy levels 290

- split read policy 210

Split-brain 53

standby path attribute 289

states

- of link objects 589

statistics gathering 177

storage

- clearing 218

- ordered allocation of 204, 206

- uninitialized 218

storage accounts 142

storage attributes and volume layout 204

storage cache 569

- used by space-optimized instant snapshots 570

Storage Checkpoints 44, 362, 571, 576

- accessing 643

- administration of 641

- converting a data Storage Checkpoint to a nodata

- Storage Checkpoint with multiple Storage
- Checkpoints 648

- creating 642

- data Storage Checkpoints 578

- definition of 640

- difference between a data Storage Checkpoint
- and a nodata Storage Checkpoint 646

- freezing and thawing a file system 572

- mounting 643

- multi-volume file system support 710

- nodata Storage Checkpoints 578, 645

- operation failures 654

- pseudo device 644

- read-only Storage Checkpoints 643

- removable Storage Checkpoints 579

- removing 643

- space management 654

- synchronous vs. asynchronous conversion 645

- types of 578

- unmounting 645

- using the fsck command 645

- writable Storage Checkpoints 643

storage disconnectivity

- CVM handling 152

storage failures 481

storage groups 142

storage processor 171

storage relayout 121

Storage Rollback 576

stripe columns 109

- stripe units
 - changing size 835
 - defined 109
- stripe-mirror-col-split-trigger-pt 202
- striped plexes
 - defined 108
- striped volumes
 - changing number of columns 835
 - changing stripe unit size 835
 - creating 201
 - defined 198
 - failure of 108
 - specifying non-default number of columns 201
 - specifying non-default stripe unit size 201
- striped-mirror volumes
 - benefits of 113
 - creating 202
 - defined 199
 - mirroring columns 202
 - mirroring subdisks 202
 - trigger point for mirroring 202
- striping 108
- striping plus mirroring 112
- subdisk names 100
- subdisks
 - blocks 100
 - complete failure messages 776
 - defined 100
 - determining failed 776
 - DRL log 128
 - hot-relocation 127, 771, 778
 - hot-relocation messages 783
 - listing original disks after hot-relocation 785
 - maximum number per plex 896
 - mirroring in striped-mirror volumes 202
 - moving after hot-relocation 783
 - partial failure messages 775
 - RAID-5 failure of 773
 - specifying different offsets for unrelocation 785
 - unrelocating after hot-relocation 783
 - unrelocating to different disks 784
 - unrelocating using vxunreloc 784
- super-block 669
- SVID requirement
 - VxFS conformance to 91
- Switching the CVM master 373
- Symmetric architecture 41
- synchronization
 - controlling for instant snapshots 622

- synchronization (*continued*)
 - improving performance of 623
- Synchronize time
 - SFCFSHA 48
- synchronous I/O 487
- syncing attribute 595, 622
- syncpause 622
- syncresume 622
- syncstart 622
- syncstop 622
- syncwait 622
- system failure recovery 78
- system performance
 - overall 213

T

- t# 94
- tags
 - for tasks 828
 - removing from volumes 843
 - renaming 843
 - setting on volumes 843
 - specifying for online relay layout tasks 836
 - specifying for tasks 828
- target IDs
 - number 94
 - specifying to vxassist 204
- target mirroring 206
- targets
 - listing 252
- task monitor in VxVM 827
- tasks
 - aborting 829
 - changing state of 829
 - identifiers 828
 - listing 829
 - managing 828
 - modifying parameters of 830
 - monitoring 829
 - monitoring online relay layout 836
 - pausing 829
 - resuming 829
 - specifying tags 828
 - specifying tags on online relay layout operation 836
 - tags 828
- temporary area used by online relay layout 122
- temporary directories 81
- terminology
 - primary and secondary 47

- Thaw 44
- thaw 489
- thin provisioning
 - using 679
- Thin Reclamation 86, 233, 685
- thin storage
 - using 679
- third-mirror
 - snapshots 131
- third-mirror break-off snapshots
 - creating 606
- third-party driver (TPD) 248
- throttling 177
- Time synchronization
 - cluster file systems 360
- tmplog mount option 216
- TPD
 - displaying path information 279
 - support for coexistence 248
- tpdmode attribute 331
- trigger
 - postoffline 512
 - postonline 512
 - preonline 511
- trigger point in striped-mirror volumes 202
- tunables
 - changing values of 906, 914
 - displaying using prtconf 922
 - dmp_cache_open 885
 - dmp_daemon_count 885
 - dmp_delayq_interval 885
 - dmp_fast_recovery 885
 - dmp_health_time 886
 - dmp_log_level 886
 - dmp_low_impact_probe 886
 - dmp_lun_retry_timeout 887
 - dmp_monitor_fabric 887
 - dmp_monitor_osevent 888
 - dmp_monitor_ownership 888
 - dmp_native_multipathing 888
 - dmp_native_support 889
 - dmp_path_age 889
 - dmp_pathswitch_blks_shift 889
 - dmp_probe_idle_lun 890
 - dmp_probe_threshold 890
 - dmp_restore_cycles 890
 - dmp_restore_interval 890
 - dmp_restore_state 891
 - dmp_scsi_timeout 891

- tunables (*continued*)
 - dmp_sfg_threshold 892
 - dmp_stat_interval 892
 - FastrResync/cache object metadata cache
 - size 903
 - maximum vol_stats_enable 895
 - vol_checkpoint_default 893
 - vol_default_iodelay 893
 - vol_fmr_logsz 135, 899
 - vol_kmsg_resend_period 893
 - vol_kmsg_send_period 893
 - vol_max_vol 894
 - vol_maxio 894
 - vol_maxioctl 894
 - vol_maxparallelio 895
 - vol_maxspecialio 895
 - vol_subdisk_num 896
 - volcvm_smartsync 904
 - voldrl_max_drtregs 900
 - voldrl_max_seq_dirty 128, 900
 - voldrl_min_regionsz 900
 - voliomem_chunk_size 896
 - voliomem_maxpool_sz 896
 - voliot_errbuf_dflt 897
 - voliot_iobuf_default 897
 - voliot_iobuf_limit 897
 - voliot_iobuf_max 897
 - voliot_max_open 898
 - volpagemod_max_memsz 903
 - volraid_minpool_size 898
 - volraid_rsrtmax 898
 - VVR
 - changing. See using vxio.conf file

- Tuning DMP
 - using templates 907

U

- umount command 225
- Unconfigure
 - Clustered NFS 515
- Understanding
 - Clustered NFS 511
- unencapsulating the root disk 883
- uninitialized storage, clearing 218
- unmount 645
 - a snapped file system 583
- Unsharing
 - shared CFS file system 519

- unsharing
 - shared CFS file system 547
- upgrade
 - from raw devices 502
- usage
 - SFCFSA 40
- use_all_paths attribute 295
- use_avid
 - vxddladm option 328
- user-specified device names 329
- utility
 - vxtune 917, 919

V

- V-5-1-2829 844
- V-5-1-569 168
- V-5-2-3091 819
- V-5-2-4292 819
- Veritas Operations Manager 35
- version 0
 - of DCOs 138
- Version 1 disk layout 939
- Version 2 disk layout 939
- version 20
 - of DCOs 138–139
- Version 3 disk layout 939
- Version 4 disk layout 940
- Version 5 disk layout 940
- Version 6 disk layout 940
- Version 7 disk layout 940
- Version 8 disk layout 940
- Version 9 disk layout 940
- versioning
 - of DCOs 138
- versions
 - disk group 848
 - displaying for disk group 849
 - upgrading 848
- virtual disks 90
- virtual objects 96
- VM disks
 - defined 99
 - determining if shared 380
 - displaying spare 779
 - excluding free space from hot-relocation use 781
 - making free space available for hot-relocation
 - use 782
 - marking as spare 779
 - mirroring volumes on 839
- VM disks *(continued)*
 - moving volumes from 812
 - names 100
 - postponing replacement 858
 - removing from pool of hot-relocation spares 780
 - renaming 348
- vol## 102
- vol###-## 102
- vol_checkpoint_default tunable 893
- vol_default_iodelay tunable 893
- vol_fmr_logsz tunable 135, 899
- vol_kmsg_resend_period tunable 893
- vol_kmsg_send_period tunable 893
- vol_max_vol tunable 894
- vol_maxio tunable 894
- vol_maxioctl tunable 894
- vol_maxparallelio tunable 895
- vol_maxspecialio tunable 895
- vol_subdisk_num tunable 896
- volbrk snapshot type 621
- volcvm_smartsync tunable 904
- voldrl_max_drtregs tunable 900
- voldrl_max_seq_dirty tunable 128, 900
- voldrl_min_regionsz tunable 900
- voliomem_chunk_size tunable 896
- voliomem_maxpool_sz tunable 896
- voliot_errbuf_dflt tunable 897
- voliot_iobuf_default tunable 897
- voliot_iobuf_limit tunable 897
- voliot_iobuf_max tunable 897
- voliot_max_open tunable 898
- volpagemod_max_memsz tunable 903
- volraid_minpool_size tunable 898
- volraid_rsrtransmax tunable 898
- volume resynchronization 126
- volume sets
 - adding volumes to 703
 - administering 701
 - controlling access to raw device nodes 707
 - creating 702
 - creating instant snapshots of 612
 - displaying access to raw device nodes 707
 - enabling access to raw device nodes 706
 - listing details of 703
 - raw device nodes 705
 - removing volumes from 703
 - starting 704
 - stopping 704

volumes

- adding mirrors 838
- adding snapshot mirrors to 614
- adding to volume sets 703
- adding version 0 DCOs to 635
- backing up 567
- backing up online using snapshots 595
- boot-time restrictions 873
- booting root 873
- changing layout online 831
- changing number of columns 835
- changing read policies for mirrored 209
- changing stripe unit size 835
- concatenated 106, 198
- concatenated-mirror 114, 199
- configuring exclusive open by cluster node 385
- configuring on SAN storage 142
- configuring site consistency on 476
- creating concatenated-mirror 201
- creating for use as full-sized instant snapshots 600
- creating from snapshots 633
- creating mirrored 200
- creating mirrored-concatenated 200
- creating mirrored-stripe 202
- creating RAID-5 203
- creating snapshots 629
- creating striped 201
- creating striped-mirror 202
- defined 102
- displaying information about snapshots 634
- dissociating version 0 DCOs from 638
- effect of growing on FastResync maps 140
- excluding storage from use by vxassist 205
- flagged as dirty 126
- layered 113, 119, 199
- limit on number of plexes 102
- limitations 102
- listing on boot (root) disk 879
- maximum number of 894
- merging snapshots 632
- mirrored 112, 198
- mirrored-concatenated 112
- mirrored-stripe 112, 199
- mirroring across controllers 208
- mirroring across targets 206
- mirroring all 838
- mirroring on disks 839
- moving from VM disks 812

volumes *(continued)*

- names 102
- naming snap 594
- performing online relayout 831
- RAID-0 108
- RAID-0+1 112
- RAID-1 112
- RAID-1+0 113
- RAID-10 113
- RAID-5 115, 199
- reattaching plexes 852
- reattaching version 0 DCOs to 638
- reconfiguration in clusters 161
- recovering after correctable hardware failure 776
- removing 854
- removing from /etc/vfstab 854
- removing linked snapshots from 615
- removing mirrors from 842
- removing plexes from 842
- removing snapshot mirrors from 615
- removing version 0 DCOs from 638
- restarting moved 826
- restoring from instant snapshots 618
- resynchronizing from snapshots 632
- snapshots 130
- spanned 106
- specifying non-default number of columns 201
- specifying non-default relayout 835
- specifying non-default stripe unit size 201
- specifying storage for version 0 DCO plexes 637
- specifying use of storage to vxassist 204
- stopping activity on 854
- striped 108, 198
- striped-mirror 113, 199
- taking multiple snapshots 594
- trigger point for mirroring in striped-mirror 202
- types of layout 198
- VX_DSYNC 487
- VX_FREEZE 489, 926
- VX_GETCACHE 489
- VX_SETCACHE 489
- VX_SNAPREAD 668
- VX_THAW 489
- VX_UNBUFFERED 487
- vxassist
 - adding DCOs to volumes 636
 - adding mirrors to volumes 838
 - configuring exclusive access to a volume 385
 - configuring site consistency on volumes 476

vxassist (continued)

- creating cache volumes 598
- creating concatenated-mirror volumes 201
- creating mirrored volumes 200
- creating mirrored-concatenated volumes 200
- creating mirrored-stripe volumes 202
- creating RAID-5 volumes 203
- creating snapshots 627
- creating striped volumes 201
- creating striped-mirror volumes 202
- creating volumes for use as full-sized instant snapshots 601
- defaults file 190
- defining layout on specified storage 204
- displaying information about snapshots 634
- dissociating snapshots from volumes 633
- excluding storage from use 205
- listing tags set on volumes 843
- merging snapshots with volumes 632
- mirroring across controllers 208
- mirroring across targets 206, 208
- mirroring file systems on root disk 878
- moving DCO plexes 637
- relaying out volumes online 831
- removing mirrors 842
- removing plexes 842
- removing tags from volumes 843
- removing version 0 DCOs from volumes 638
- removing volumes 854
- replacing tags set on volumes 843
- resynchronizing volumes from snapshots 632
- setting default values 190
- setting tags on volumes 843–844
- snapabort 586
- snapback 587
- snapshot 587
- snapstart 586
- specifying number of mirrors 200
- specifying number of RAID-5 logs 203
- specifying ordered allocation of storage 206
- specifying plexes for online relayout 836
- specifying storage attributes 204
- specifying storage for version 0 DCO plexes 637
- specifying tags for online relayout tasks 836
- taking snapshots of multiple volumes 631

vxcache

- listing snapshots in a cache 624
- resizing caches 625
- starting cache objects 599

vxcache (continued)

- stopping a cache 626
- tuning cache autogrow 625

vxcached

- tuning 624

vxclustadm 159*vxconfigd*

- managing with vxctl 103
- monitoring configuration changes 831
- operation in clusters 162

vxdaresore

- handling simple/nopriv disk failures 332

vxdc

- dissociating version 0 DCOs from volumes 638
- reattaching version 0 DCOs to volumes 638
- removing version 0 DCOs from volumes 638

vxctl

- checking cluster protocol version 386
- managing vxconfigd 103
- setting a site tag 471, 474
- setting default disk group 192
- usage in clusters 379

vxctl enable

- invoking device discovery 247

vxddladm

- adding disks to DISKS category 259
- adding foreign devices 261
- changing naming scheme 328
- displaying the disk-naming scheme 328
- listing all devices 251
- listing configured devices 254
- listing configured targets 253
- listing excluded disk arrays 257, 259
- listing ports on a Host Bus Adapter 252
- listing supported disk arrays 256
- listing supported disks in DISKS category 257
- listing supported HBAs 251
- removing disks from DISKS category 249, 260–261
- setting iSCSI parameters 254
- used to exclude support for disk arrays 257
- used to re-include support for disk arrays 257

vxdg

- changing activation mode on shared disk groups 384
- configuring site consistency for a disk group 472
- configuring site-based allocation for a disk group 471
- converting shared disk groups to private 383

vx dg (continued)

- creating disk groups with old version number 849
- creating shared disk groups 381
- displaying boot disk group 192
- displaying default disk group 192
- displaying disk group version 849
- importing shared disk groups 382
- joining disk groups 825
- listing objects affected by move 819
- listing shared disk groups 380
- listing spare disks 779
- moving disks between disk groups 813
- moving objects between disk groups 821
- reattaching a site 473
- setting a site name 476, 479
- setting disk group policies 384
- simulating site failure 473
- splitting disk groups 824
- upgrading disk group version 849

vx disk

- determining if disks are shared 380
- displaying multi-pathing information 268
- listing disks 326
- listing spare disks 779
- notifying dynamic LUN expansion 318
- scanning disk devices 244
- setting a site name 475

vx disk scandisks

- rescanning devices 244
- scanning devices 244

vx diskadd

- placing disks under VxVM control 345

vx diskadm

- Add or initialize one or more disks 336
- adding disks 336
- changing the disk-naming scheme 327
- Encapsulate one or more disks 868
- Exclude a disk from hot-relocation use 781
- excluding free space on disks from hot-relocation use 781
- initializing disks 336
- List disk information 326
- listing spare disks 779
- Make a disk available for hot-relocation use 782
- making free space on disks available for hot-relocation use 782
- Mark a disk as a spare for a disk group 780
- marking disks as spare 780
- Mirror volumes on a disk 839

vx diskadm (continued)

- mirroring disks 873
- mirroring root disks 876
- mirroring volumes 839
- Move volumes from a disk 813
- moving disks between disk groups 814
- moving volumes from VM disks 813
- Remove a disk 346
- Remove a disk for replacement 859
- removing disks from pool of hot-relocation spares 781
- Replace a failed or removed disk 862
- Turn off the spare flag on a disk 781

vx diskconfig

- purpose of 243

vx diskunsetup

- removing disks from VxVM control 855

vx dmpadm

- changing TPD naming scheme 331
- configuring an APM 308
- configuring I/O throttling 302
- configuring response to I/O errors 301, 304
- disabling controllers in DMP 265
- disabling I/O in DMP 298
- displaying APM information 308
- displaying DMP database information 265
- displaying DMP node for a path 271, 273
- displaying DMP node for an enclosure 271–272
- displaying I/O error recovery settings 304
- displaying I/O policy 291
- displaying I/O throttling settings 304
- displaying information about controllers 276
- displaying information about enclosures 277
- displaying partition size 291
- displaying paths controlled by DMP node 274
- displaying status of DMP restoration thread 307
- displaying TPD information 279
- enabling I/O in DMP 299
- gathering I/O statistics 283
- listing information about array ports 278
- removing an APM 308
- renaming enclosures 300
- setting I/O policy 294–295
- setting path attributes 289
- setting restore polling interval 305
- specifying DMP path restoration policy 305
- stopping DMP restore daemon 307

vx dmpadm list

- displaying DMP nodes 272

- vxdump 239
- vxedit
 - excluding free space on disks from hot-relocation use 781
 - making free space on disks available for hot-relocation use 782
 - marking disks as spare 779
 - removing a cache 626
 - removing disks from pool of hot-relocation spares 780
 - removing instant snapshots 619
 - removing snapshots from a cache 626
 - removing volumes 854
 - renaming disks 348
- VxFS
 - storage allocation 212
 - supported features
 - cluster file systems 42
 - unsupported features
 - cluster file systems 45
- vxfs_inotopath 937
- vxiod I/O kernel threads 93
- vxlsino 937
- vxmake
 - creating cache objects 598
 - creating plexes 838
- vxmend
 - re-enabling plexes 852
- vxmirror
 - configuring VxVM default behavior 838
 - mirroring root disks 876
 - mirroring volumes 838
- vxnotify
 - monitoring configuration changes 831
- vxplex
 - attaching plexes to volumes 838
 - converting plexes to snapshots 630
 - reattaching plexes 852
 - removing mirrors 842
 - removing mirrors of root disk volumes 883
 - removing plexes 842
- vxprint
 - displaying DCO information 637
 - displaying snapshots configured on a cache 626
 - listing spare disks 779
 - listing volumes on boot disks 879
 - verifying if volumes are prepared for instant snapshots 597
- vxrecover
 - recovering plexes 776
 - restarting moved volumes 826
- vxrelayout
 - resuming online relayout 837
 - reversing direction of online relayout 837
 - viewing status of online relayout 836
- vxrelocd
 - hot-relocation daemon 772
 - modifying behavior of 786
 - notifying users other than root 787
 - operation of 773
 - preventing from running 787
 - reducing performance impact of recovery 787
- vxrestore 239
- vxsnap
 - adding snapshot mirrors to volumes 614
 - administering instant snapshots 588
 - backing up multiple volumes 611
 - controlling instant snapshot synchronization 622
 - creating a cascaded snapshot hierarchy 615
 - creating full-sized instant snapshots 604, 610
 - creating linked break-off snapshot volumes 610
 - creating space-optimized instant snapshots 602
 - displaying information about instant snapshots 620
 - dissociating instant snapshots 618
 - preparing volumes for instant snapshots 597
 - reattaching instant snapshots 616
 - reattaching linked third-mirror snapshots 617
 - refreshing instant snapshots 616
 - removing a snapshot mirror from a volume 615
 - restore 588
 - restoring volumes 618
 - splitting snapshot hierarchies 619
- vxspcshow
 - discovering device names 143
- vxstat
 - determining which disks have failed 776
 - usage with clusters 387
- vxtask
 - aborting tasks 830
 - listing tasks 830
 - monitoring online relayout 836
 - monitoring tasks 830
 - pausing online relayout 837
 - resuming online relayout 837
 - resuming tasks 830

- vxtune
 - setting volpagemod_max_memsz 903
- vxtune utility 917, 919
- vxunrelloc
 - listing original disks of hot-relocated subdisks 785
 - moving subdisks after hot-relocation 784
 - restarting after errors 786
 - specifying different offsets for unrelocated subdisks 785
 - unrelocating subdisks after hot-relocation 784
 - unrelocating subdisks to different disks 784
- vxunroot
 - removing rootability 883
 - unencapsulating the root disk 883
- VxVM
 - cluster functionality (CVM) 145
 - configuration daemon 103
 - configuring to create mirrored volumes 838
 - dependency on operating system 92
 - disk discovery 245
 - granularity of memory allocation by 896
 - limitations of shared disk groups 157
 - maximum number of subdisks per plex 896
 - maximum number of volumes 894
 - maximum size of memory pool 896
 - minimum size of memory pool 898
 - objects in 96
 - operation in clusters 147
 - removing disks from control of 855
 - rootability 872
 - shared objects in cluster 150
 - task monitor 827
 - types of volume layout 198
 - upgrading 848
 - upgrading disk group version 849
- vxvol
 - configuring exclusive access to a volume 385
 - configuring site consistency on volumes 476
 - restarting moved volumes 826
 - setting read policy 210
 - stopping volumes 854
- vxvset
 - adding volumes to volume sets 703
 - controlling access to raw device nodes 707
 - creating volume sets 702
 - creating volume sets with raw device access 706
 - listing details of volume sets 703
 - removing volumes from volume sets 703
 - starting volume sets 704

- vxvset (*continued*)
 - stopping volume sets 704

W

- warning messages
 - Specified region-size is larger than the limit on the system 596
- worldwide name identifiers 327
- writable Storage Checkpoints 643
- write size 237
- WWN identifiers 327