

Veritas™ Storage Foundation 6.0.1 Administrator's Guide - Solaris

Veritas™ Storage Foundation 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: 6.0.1

Document version: 6.0.1 Rev 2

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	25
Chapter 1 Overview of Storage Foundation	26
About Veritas Storage Foundation	26
About Veritas File System	27
About the Veritas File System intent log	27
About extents	28
About file system disk layouts	29
About Veritas Volume Manager	29
About Veritas Dynamic Multi-Pathing (DMP)	30
About Veritas Operations Manager	30
About Storage Foundation solutions	31
Chapter 2 How Veritas File System works	32
Veritas File System features	32
Veritas File System performance enhancements	41
Enhanced I/O performance	42
Delayed allocation for extending writes	43
Using Veritas File System	44
Online system administration	44
Application program interface	45
Chapter 3 How Veritas Volume Manager works	47
How Veritas Volume Manager works with the operating system	47
How data is stored	48
How Veritas Volume Manager handles storage management	49
Physical objects	49
Virtual objects	51
About the configuration daemon in Veritas Volume Manager	58
Multiple paths to disk arrays	59
Volume layouts in Veritas Volume Manager	59
Non-layered volumes	60

Layered volumes	60
Layout methods	60
Concatenation, spanning, and carving	61
Striping (RAID-0)	63
Mirroring (RAID-1)	66
Striping plus mirroring (mirrored-stripe or RAID-0+1)	67
Mirroring plus striping (striped-mirror, RAID-1+0, or RAID-10)	68
RAID-5 (striping with parity)	69
Online relay layout	76
How online relay layout works	76
Limitations of online relay layout	79
Transformation characteristics	80
Transformations and volume length	80
Volume resynchronization	81
Dirty flags	81
Resynchronization process	81
Hot-relocation	82
Dirty region logging	82
Log subdisks and plexes	83
Sequential DRL	83
SmartSync recovery accelerator	83
Volume snapshots	85
Comparison of snapshot features	86
FastResync	87
How FastResync works	88
How non-persistent FastResync works with snapshots	89
How persistent FastResync works with snapshots	90
DCO volume versioning	93
Effect of growing a volume on the FastResync map	95
FastResync limitations	96
Volume sets	97
Configuration of volumes on SAN storage	97

Chapter 4	How Veritas Dynamic Multi-Pathing works	100
	How DMP works	100
	Device discovery	103
	How DMP monitors I/O on paths	106
	Load balancing	108
	Disabling Sun Multipath IO (MPxIO)	109
	Dynamic Reconfiguration	109
	About booting from DMP devices	110

	DMP in a clustered environment	110
	Veritas Volume Manager co-existence with Oracle Automatic Storage Management disks	111
Section 2	Provisioning storage	114
Chapter 5	Provisioning new storage	115
	Provisioning new storage	115
	Growing the existing storage by adding a new LUN	116
	Growing the existing storage by growing the LUN	116
	Displaying SF information with vxlist	117
Chapter 6	Advanced allocation methods for configuring storage	118
	Customizing allocation behavior	118
	Setting default values for vxassist	120
	Using rules to make volume allocation more efficient	121
	Understanding persistent attributes	124
	Customizing disk classes for allocation	127
	Specifying allocation constraints for vxassist operations with the use clause and the require clause	129
	Management of the use and require type of persistent attributes	137
	Creating volumes of a specific layout	141
	Types of volume layouts	142
	Creating a mirrored volume	143
	Creating a striped volume	145
	Creating a RAID-5 volume	146
	Creating a volume on specific disks	148
	Creating volumes on specific media types	150
	Specifying ordered allocation of storage to volumes	150
	Site-based allocation	153
	Changing the read policy for mirrored volumes	153
Chapter 7	Creating and mounting VxFS file systems	155
	Creating a VxFS file system	155
	Block size	156
	Intent log size	157
	Example of creating a file system	157
	Converting a file system to VxFS	158
	Example of converting a file system	159

Mounting a VxFS file system	159
log mount option	161
delaylog mount option	161
tmplog mount option	162
logiosize mount option	163
nodatainlog mount option	163
blkclear mount option	163
mincache mount option	163
convosync mount option	165
ioerror mount option	166
largefiles and nolargefiles mount options	167
cio mount option	169
mntlock mount option	169
ckptautomnt mount option	169
Combining mount command options	169
Example of mounting a file system	170
Unmounting a file system	170
Example of unmounting a file system	171
Resizing a file system	171
Extending a file system using fsadm	171
Shrinking a file system	172
Reorganizing a file system	173
Displaying information on mounted file systems	175
Example of displaying information on mounted file systems	175
Identifying file system types	175
Example of determining a file system's type	176
Monitoring free space	176
Monitoring fragmentation	177

Chapter 8	Extent attributes	179
	About extent attributes	179
	Reservation: preallocating space to a file	180
	Fixed extent size	180
	How the fixed extent size works with the shared extents	181
	Other extent attribute controls	181
	Commands related to extent attributes	183
	Example of setting an extent attribute	184
	Example of getting an extent attribute	184
	Failure to preserve extent attributes	184

Section 3	Administering multi-pathing with DMP	186
Chapter 9	Administering Dynamic Multi-Pathing	187
	Discovering and configuring newly added disk devices	187
	Partial device discovery	188
	Discovering disks and dynamically adding disk arrays	189
	Third-party driver coexistence	191
	How to administer the Device Discovery Layer	193
	Making devices invisible to VxVM	205
	Making devices visible to VxVM	207
	About enabling and disabling I/O for controllers and storage processors	207
	About displaying DMP database information	208
	Displaying the paths to a disk	208
	Administering DMP using vxddmpadm	211
	Retrieving information about a DMP node	213
	Displaying consolidated information about the DMP nodes	214
	Displaying the members of a LUN group	215
	Displaying paths controlled by a DMP node, controller, enclosure, or array port	215
	Displaying information about controllers	218
	Displaying information about enclosures	219
	Displaying information about array ports	220
	Displaying information about TPD-controlled devices	220
	Displaying extended device attributes	222
	Suppressing or including devices from VxVM control	224
	Gathering and displaying I/O statistics	225
	Setting the attributes of the paths to an enclosure	230
	Displaying the redundancy level of a device or enclosure	232
	Specifying the minimum number of active paths	233
	Displaying the I/O policy	233
	Specifying the I/O policy	234
	Disabling I/O for paths, controllers, array ports, or DMP nodes	240
	Enabling I/O for paths, controllers, array ports, or DMP nodes	241
	Renaming an enclosure	242
	Configuring the response to I/O failures	243
	Configuring the I/O throttling mechanism	244
	Configuring Low Impact Path Probing	245

	Configuring Subpaths Failover Groups (SFG)	246
	Displaying recovery option values	246
	Configuring DMP path restoration policies	247
	Stopping the DMP path restoration thread	249
	Displaying the status of the DMP path restoration thread	249
	Configuring array policy modules	249
	DMP coexistence with native multi-pathing	250
Chapter 10	Dynamic Reconfiguration of devices	252
	About online Dynamic Reconfiguration	252
	Reconfiguring a LUN online that is under DMP control	253
	Removing LUNs dynamically from an existing target ID	253
	Adding new LUNs dynamically to a new target ID	255
	Replacing LUNs dynamically from an existing target ID	256
	Dynamic LUN expansion	257
	Changing the characteristics of a LUN from the array side	259
	Replacing a host bus adapter online	260
	Upgrading the array controller firmware online	260
Chapter 11	Managing devices	262
	Displaying disk information	262
	Displaying disk information with vxdiskadm	263
	Changing the disk device naming scheme	264
	Displaying the disk-naming scheme	265
	Setting customized names for DMP nodes	266
	Regenerating persistent device names	267
	Changing device naming for TPD-controlled enclosures	267
	Simple or nopriv disks with enclosure-based naming	269
	About the Array Volume Identifier (AVID) attribute	270
	About disk installation and formatting	272
	Adding and removing disks	273
	Adding a disk to VxVM	273
	Removing disks	282
	Renaming a disk	285
Chapter 12	Event monitoring	287
	About the Dynamic Multi-Pathing (DMP) event source daemon (vxesd)	287
	Fabric Monitoring and proactive error detection	288
	Dynamic Multi-Pathing (DMP) automated device discovery	289

	Dynamic Multi-Pathing (DMP) discovery of iSCSI and SAN Fibre	
	Channel topology	290
	DMP event logging	290
	Starting and stopping the Dynamic Multi-Pathing (DMP) event source	
	daemon	291
Section 4	Optimizing I/O performance	292
Chapter 13	Veritas File System I/O	293
	About Veritas File System I/O	293
	Buffered and Direct I/O	294
	Direct I/O	294
	Unbuffered I/O	295
	Data synchronous I/O	295
	Concurrent I/O	296
	Cache advisories	297
	Freezing and thawing a file system	297
	Getting the I/O size	298
	About Storage Foundation and High Availability Solutions products	
	database accelerators	298
Chapter 14	Veritas Volume Manager I/O	300
	Veritas Volume Manager throttling of administrative I/O	300
Section 5	Using Point-in-time copies	302
Chapter 15	Understanding point-in-time copy methods	303
	About point-in-time copies	303
	When to use point-in-time copies	304
	Implementing point-in time copy solutions on a primary host	305
	Implementing off-host point-in-time copy solutions	307
	About Storage Foundation point-in-time copy technologies	313
	Comparison of Point-in-time copy solutions	314
	Volume-level snapshots	315
	Persistent FastResync of volume snapshots	315
	Data integrity in volume snapshots	315
	Third-mirror break-off snapshots	316
	Space-optimized instant volume snapshots	317
	Choices for snapshot resynchronization	318
	Disk group split/join	318

Storage Checkpoints	319
How Storage Checkpoints differ from snapshots	320
How a Storage Checkpoint works	320
Types of Storage Checkpoints	324
About FileSnaps	327
Properties of FileSnaps	327
Concurrent I/O to FileSnaps	328
Copy-on-write and FileSnaps	328
Reading from FileSnaps	329
Block map fragmentation and FileSnaps	329
Backup and FileSnaps	329
About snapshot file systems	330
How a snapshot file system works	330

Chapter 16	Administering volume snapshots	332
	About volume snapshots	332
	How traditional third-mirror break-off snapshots work	333
	How full-sized instant snapshots work	334
	Linked break-off snapshot volumes	336
	Cascaded snapshots	337
	Creating a snapshot of a snapshot	338
	Creating multiple snapshots	340
	Restoring the original volume from a snapshot	341
	Creating instant snapshots	342
	Adding an instant snap DCO and DCO volume	343
	Creating and managing space-optimized instant snapshots	349
	Creating and managing full-sized instant snapshots	352
	Creating and managing third-mirror break-off snapshots	354
	Creating and managing linked break-off snapshot volumes	357
	Creating multiple instant snapshots	359
	Creating instant snapshots of volume sets	360
	Adding snapshot mirrors to a volume	362
	Removing a snapshot mirror	363
	Removing a linked break-off snapshot volume	363
	Adding a snapshot to a cascaded snapshot hierarchy	363
	Refreshing an instant space-optimized snapshot	364
	Reattaching an instant full-sized or plex break-off snapshot	364
	Reattaching a linked break-off snapshot volume	365
	Restoring a volume from an instant space-optimized snapshot	366
	Dissociating an instant snapshot	366
	Removing an instant snapshot	367

Splitting an instant snapshot hierarchy	367
Displaying instant snapshot information	368
Controlling instant snapshot synchronization	370
Listing the snapshots created on a cache	372
Tuning the autogrow attributes of a cache	372
Monitoring and displaying cache usage	373
Growing and shrinking a cache	373
Removing a cache	374
Creating traditional third-mirror break-off snapshots	374
Converting a plex into a snapshot plex	378
Creating multiple snapshots with the vxassist command	379
Reattaching a snapshot volume	380
Adding plexes to a snapshot volume	381
Dissociating a snapshot volume	381
Displaying snapshot information	382
Adding a version 0 DCO and DCO volume	383
Specifying storage for version 0 DCO plexes	384
Removing a version 0 DCO and DCO volume	386
Reattaching a version 0 DCO and DCO volume	386

Chapter 17 Administering Storage Checkpoints 388

About Storage Checkpoints	388
Storage Checkpoint administration	389
Creating a Storage Checkpoint	390
Removing a Storage Checkpoint	391
Accessing a Storage Checkpoint	391
Converting a data Storage Checkpoint to a nodata Storage Checkpoint	393
Enabling and disabling Storage Checkpoint visibility	401
Storage Checkpoint space management considerations	402
Restoring from a Storage Checkpoint	402
Examples of restoring a file from a Storage Checkpoint	403
Storage Checkpoint quotas	408

Chapter 18 Administering FileSnaps 409

FileSnap creation	409
FileSnap creation over Network File System	409
Using FileSnaps	410
Using FileSnaps to create point-in-time copies of files	411
Using FileSnaps to provision virtual desktops	411
Using FileSnaps to optimize write intensive applications for virtual machines	412

	Using FileSnaps to create multiple copies of data instantly	412
	Comparison of the logical size output of the fsadm -S shared, du, and df commands	413
Chapter 19	Administering snapshot file systems	415
	Snapshot file system backups	415
	Snapshot file system performance	416
	About snapshot file system disk structure	416
	Differences between snapshots and Storage Checkpoints	417
	Creating a snapshot file system	418
	Examples of creating snapshot file systems	419
Section 6	Optimizing storage with Storage Foundation	420
Chapter 20	Understanding storage optimization solutions in Storage Foundation	421
	About thin provisioning	421
	About thin optimization solutions in Storage Foundation	422
	About SmartMove	423
	SmartMove for thin provisioning	423
	About the Thin Reclamation feature	424
	Determining when to reclaim space on a thin reclamation LUN	424
	How automatic reclamation works	425
Chapter 21	Migrating data from thick storage to thin storage	427
	About using SmartMove to migrate to Thin Storage	427
	Migrating to thin provisioning	427
Chapter 22	Maintaining Thin Storage with Thin Reclamation	431
	Reclamation of storage on thin reclamation arrays	431
	About Thin Reclamation of a disk, a disk group, or an enclosure	432
	About Thin Reclamation of a file system	433
	Identifying thin and thin reclamation LUNs	433
	Displaying VxFS file system usage on thin reclamation LUNs	434
	Reclaiming space on a file system	437

	Reclaiming space on a disk, disk group, or enclosure	438
	About the reclamation log file	439
	Monitoring Thin Reclamation using the vxtask command	440
	Configuring automatic reclamation	442
Section 7	Maximizing storage utilization	444
Chapter 23	Understanding storage tiering with SmartTier	445
	About SmartTier	445
	About VxFS multi-volume file systems	447
	About VxVM volume sets	448
	About volume tags	448
	SmartTier file management	448
	SmartTier sub-file object management	449
	How the SmartTier policy works with the shared extents	449
	SmartTier in a High Availability (HA) environment	450
Chapter 24	Creating and administering volume sets	451
	About volume sets	451
	Creating a volume set	452
	Adding a volume to a volume set	453
	Removing a volume from a volume set	453
	Listing details of volume sets	453
	Stopping and starting volume sets	454
	Managing raw device nodes of component volumes	455
	Enabling raw device access when creating a volume set	456
	Displaying the raw device access settings for a volume set	457
	Controlling raw device access for an existing volume set	457
Chapter 25	Multi-volume file systems	459
	About multi-volume file systems	459
	About volume types	460
	Features implemented using multi-volume support	460
	Volume availability	461
	Creating multi-volume file systems	462
	Example of creating a multi-volume file system	462
	Converting a single volume file system to a multi-volume file system	463
	Adding a volume to and removing a volume from a multi-volume file system	465
	Adding a volume to a multi-volume file system	465

Removing a volume from a multi-volume file system	465
Forcibly removing a volume in a multi-volume file system	465
Moving volume 0 in a multi-volume file system	466
Volume encapsulation	466
Encapsulating a volume	466
Deencapsulating a volume	468
Reporting file extents	468
Examples of reporting file extents	469
Load balancing	470
Defining and assigning a load balancing allocation policy	471
Rebalancing extents	471
Converting a multi-volume file system to a single volume file system	472

Chapter 26	Administering SmartTier	474
	About SmartTier	474
	About compressing files with SmartTier	476
	Supported SmartTier document type definitions	476
	Placement classes	477
	Tagging volumes as placement classes	478
	Listing placement classes	478
	Administering placement policies	479
	Assigning a placement policy	479
	Unassigning a placement policy	480
	Analyzing the space impact of enforcing a placement policy	480
	Querying which files will be affected by enforcing a placement policy	480
	Enforcing a placement policy	480
	Validating a placement policy	482
	File placement policy grammar	482
	File placement policy rules	483
	SELECT statement	483
	CREATE statement	486
	RELOCATE statement	488
	DELETE statement	503
	COMPRESS statement	505
	UNCOMPRESS statement	515
	Calculating I/O temperature and access temperature	524
	Multiple criteria in file placement policy rule statements	528
	Multiple file selection criteria in SELECT statement clauses	529

Multiple placement classes in <ON> clauses of CREATE statements and in <TO> clauses of RELOCATE statements	530
Multiple placement classes in <FROM> clauses of RELOCATE and DELETE statements	531
Multiple conditions in <WHEN> clauses of RELOCATE and DELETE statements	531
File placement policy rule and statement ordering	531
File placement policies and extending files	534
Using SmartTier with solid state disks	534
Fine grain temperatures with solid state disks	535
Prefer mechanism with solid state disks	535
Average I/O activity with solid state disks	536
Frequent SmartTier scans with solid state disks	536
Quick identification of cold files with solid state disks	537
Example placement policy when using solid state disks	538
Sub-file relocation	542
Moving sub-file data of files to specific target tiers	542

Chapter 27 Administering hot-relocation 543

About hot-relocation	543
How hot-relocation works	544
Partial disk failure mail messages	547
Complete disk failure mail messages	548
How space is chosen for relocation	549
Configuring a system for hot-relocation	550
Displaying spare disk information	551
Marking a disk as a hot-relocation spare	551
Removing a disk from use as a hot-relocation spare	552
Excluding a disk from hot-relocation use	553
Making a disk available for hot-relocation use	554
Configuring hot-relocation to use only spare disks	554
Moving relocated subdisks	555
Moving relocated subdisks using vxunreloc	555
Restarting vxunreloc after errors	558
Modifying the behavior of hot-relocation	558

Chapter 28 Deduplicating data on Solaris SPARC 560

About deduplicating data	560
About deduplication chunk size	561
Deduplication and file system performance	562
About the deduplication scheduler	562

	Deduplicating data	563
	Enabling and disabling deduplication on a file system	564
	Scheduling deduplication of a file system	564
	Performing a deduplication dry run	565
	Querying the deduplication status of a file system	566
	Starting and stopping the deduplication scheduler daemon	566
	Example of deduplicating a file system	567
	Deduplication results	568
	Deduplication supportability	569
	Deduplication use cases	569
	Deduplication limitations	569
Chapter 29	Compressing files	571
	About compressing files	571
	About the compressed file format	572
	About the file compression attributes	572
	About the file compression block size	573
	Compressing files with the vxcompress command	573
	Examples of using the vxcompress command	574
	Interaction of compressed files and other commands	575
	Interaction of compressed files and other features	576
	Interaction of compressed files and applications	577
	Use cases for compressing files	578
	Compressed files and databases	578
	Compressing all files that meet the specified criteria	582
Section 8	Administering storage	583
Chapter 30	Managing volumes and disk groups	584
	Rules for determining the default disk group	584
	Displaying the system-wide boot disk group	585
	Displaying and specifying the system-wide default disk group	585
	Moving volumes or disks	586
	Moving volumes from a VM disk	586
	Moving disks between disk groups	587
	Reorganizing the contents of disk groups	588
	Monitoring and controlling tasks	601
	Specifying task tags	601
	Managing tasks with vxtask	602
	Using vxnotify to monitor configuration changes	604

Performing online relayout	605
Permitted relayout transformations	606
Specifying a non-default layout	608
Specifying a plex for relayout	609
Tagging a relayout operation	609
Viewing the status of a relayout	610
Controlling the progress of a relayout	610
Adding a mirror to a volume	611
Mirroring all volumes	612
Mirroring volumes on a VM disk	612
Additional mirroring considerations	613
Configuring SmartMove	614
Removing a mirror	615
Setting tags on volumes	616
Managing disk groups	617
Disk group versions	617
Displaying disk group information	623
Creating a disk group	625
Removing a disk from a disk group	626
Deporting a disk group	627
Importing a disk group	628
Handling of minor number conflicts	630
Moving disk groups between systems	631
Handling cloned disks with duplicated identifiers	637
Renaming a disk group	646
Handling conflicting configuration copies	648
Disabling a disk group	655
Destroying a disk group	655
Backing up and restoring disk group configuration data	656
Working with existing ISP disk groups	656
Managing plexes and subdisks	658
Reattaching plexes	658
Plex synchronization	661
Decommissioning storage	662
Removing a volume	662
Removing a disk from VxVM control	663
About shredding data	663
Shredding a VxVM disk	664
Failed disk shred operation results in a disk with no label	667
Removing and replacing disks	667

Chapter 31	Rootability	675
	Encapsulating a disk	675
	Failure of disk encapsulation	679
	Using nopriv disks for encapsulation	679
	Rootability	681
	Booting root volumes	682
	Boot-time volume restrictions	682
	Creating redundancy for the root disk	683
	Creating an archived back-up root disk for disaster recovery	684
	Mirroring an encapsulated root disk	684
	Booting from alternate boot disks	685
	Mirroring other file systems on the root disk	687
	Encapsulating SAN disks	688
	Administering an encapsulated boot disk	689
	Creating a snapshot of an encapsulated boot disk	690
	Growing an encapsulated boot disk	690
	Unencapsulating the root disk	692
Chapter 32	Quotas	693
	About quota limits	693
	About quota files on Veritas File System	694
	About quota commands	694
	About quota checking with Veritas File System	696
	Using quotas	696
	Turning on quotas	696
	Turning on quotas at mount time	697
	Editing user and group quotas	697
	Modifying time limits	698
	Viewing disk quotas and usage	698
	Displaying blocks owned by users or groups	699
	Turning off quotas	699
Chapter 33	File Change Log	700
	About File Change Log	700
	About the File Change Log file	701
	File Change Log administrative interface	701
	File Change Log programmatic interface	704
	Summary of API functions	706

Section 9	Reference	707
Chapter 34	Reverse path name lookup	708
	About reverse path name lookup	708
Appendix A	Tunable parameters	710
	About tuning Veritas Storage Foundation	710
	Tuning the VxFS file system	710
	Tuning inode table size	711
	Tuning performance optimization of inode allocation	711
	Partitioned directories	711
	vx_maxlink	712
	Veritas Volume Manager maximum I/O size	712
	DMP tunable parameters	713
	Methods to change Veritas Dynamic Multi-Pathing tunable parameters	721
	Changing the values of DMP parameters with the vxddmcmd	721
	About tuning Veritas Dynamic Multi-Pathing (DMP) with templates	722
	Tunable parameters for VxVM	729
	Tunable parameters for core VxVM	730
	Tunable parameters for FlashSnap (FMR)	736
	Tunable parameters for CVM	742
	Tunable parameters for VVR	742
	Points to note when changing the values of the VVR tunables	743
	Methods to change Veritas Volume Manager tunable parameters	744
	Changing the values of the Veritas Volume Manager tunable parameters using the vxtune command line	745
	Changing the value of the Veritas Volume Manager tunable parameters using templates	747
	Changing the tunable values using the vxio.conf file	748
Appendix B	Veritas File System disk layout	750
	About Veritas File System disk layouts	750
	VxFS Version 7 disk layout	752
	VxFS Version 8 disk layout	752
	VxFS Version 9 disk layout	753

Appendix C Command reference 755

 Command completion for Veritas commands 755

 Veritas Volume Manager command reference 757

 Veritas Volume Manager manual pages 777

 Section 1M — administrative commands 777

 Section 4 — file formats 781

 Section 7 — device driver interfaces 781

 Veritas File System command summary 782

 Veritas File System manual pages 784

Index 791

Introducing Storage Foundation

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

Overview of Storage Foundation

This chapter includes the following topics:

- [About Veritas Storage Foundation](#)
- [About Veritas File System](#)
- [About Veritas Volume Manager](#)
- [About Veritas Dynamic Multi-Pathing \(DMP\)](#)
- [About Veritas Operations Manager](#)
- [About Storage Foundation solutions](#)

About Veritas Storage Foundation

Veritas Storage Foundation by Symantec includes Veritas File System (VxFS) and Veritas Volume Manager (VxVM.)

Veritas File System is a high performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.

Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.

VxFS and VxVM are included in all Veritas Storage Foundation products. If you have purchased a Veritas Storage Foundation product, VxFS and VxVM are installed

and updated as part of that product. Do not install or update them as individual components.

Veritas Storage Foundation includes the dynamic multi-pathing functionality.

The Veritas Replicator option, which replicates data to remote locations over an IP network, can also be licensed with this product.

Before you install the product, read the *Veritas Storage Foundation Release Notes*.

To install the product, follow the instructions in the *Veritas Storage Foundation Installation Guide*.

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:

File system logging

[About the Veritas File System intent log](#)

Extents

[About extents](#)

File system disk layouts

[About file system disk layouts](#)

About the Veritas File System intent log

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 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. 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.

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 might not completely recover the damaged file system structure if there was a disk hardware failure; hardware problems might 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 `fsck_vxfs(1M)` manual page and `mount_vxfs(1M)` manual page.

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. You can resize the intent log at a later time by using the `fsadm` command.

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

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 [“Intent log size”](#) on page 157.

About 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. 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). Since Veritas File System (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.

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.

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 1 KB for file system sizes of up to 1 TB, and 8 KB for file system sizes 1 TB or larger.

About file system disk layouts

The disk layout is the way file system information is stored on disk. On Veritas File System (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 (DMP)

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 Veritas 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 (SF) Enterprise license, a SF HA Enterprise license, and a Storage Foundation 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

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.

Symantec recommends using Veritas Operations Manager (VOM) to manage Storage Foundation and Cluster Server environments.

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.

About Storage Foundation solutions

Storage Foundation components and features can be used individually and together to improve performance, resilience, and ease of management for your storage and applications. Storage Foundation features can be used for:

- Improving database performance: you can use Storage Foundation 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 thin provisioning and thin reclamation solutions to set up and maintain thin storage.
- Backing up and recovering data: you can use Storage Foundation 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 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 point-in-time copy methods.
- Optimizing virtual desktop environments: you can use Storage Foundation FileSnap to optimize your virtual desktop environment.
- Maximizing storage utilization: you can use Storage Foundation SmartTier to move data to storage tiers based on age, priority, and access rate criteria.
- Migrating your data: you can use Storage Foundation Portable Data Containers to easily and reliably migrate data from one environment to another.

For a supplemental guide that documents Storage Foundation use case solutions using example scenarios: See the *Veritas Storage Foundation™ and High Availability Solutions Guide*.

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

[Table 2-1](#) lists the Veritas File System (VxFS) features.

Table 2-1 Veritas File System features

Feature	Description
Extent-based allocation	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. 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). Since Veritas File System (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. See “About extents” on page 28.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
Extent attributes	Veritas File System (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. See “About extent attributes” on page 179.
Fast file system recovery	Most file systems rely on full structural verification by the <code>fsck</code> 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. See “About the Veritas File System intent log” on page 27.
Extended mount options	The VxFS file system provides the following enhancements to the <code>mount</code> 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 159.
Enhanced data integrity modes	VxFS has the following <code>mount</code> command options to enable the enhanced data integrity modes: ■ <code>blkclear</code> See “blkclear mount option” on page 163. ■ <code>closesync</code> See “mincache mount option” on page 163. ■ <code>log</code> See “log mount option” on page 161.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
Enhanced performance mode	The default VxFS logging mode, <code>mount -o delaylog</code>, increases performance by delaying the logging of some structural changes. However, <code>delaylog</code> does not provide the equivalent data integrity as the enhanced data integrity 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. See the <code>mount_vxfs(1M)</code> manual page. See “delaylog mount option” on page 161.
Temporary file system mode	On most UNIX systems, temporary file system directories, such as <code>/tmp</code> and <code>/usr/tmp</code>, 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 <code>mount -o tmplog</code> option, which allows the user to achieve higher performance on temporary file systems by delaying the logging of most operations. See the <code>mount_vxfs(1M)</code> manual page. See “tmplog mount option” on page 162.
Improved synchronous writes	VxFS provides superior performance for synchronous write applications. The <code>mount -o datainlog</code> option greatly improves the performance of small synchronous writes. The <code>mount -o convosync=dsync</code> option improves the performance of applications that require synchronous data writes but not synchronous inode time updates. See the <code>mount_vxfs(1M)</code> manual page. Warning: The use of the <code>-o convosync=dsync</code> option violates POSIX semantics. See “convosync mount option” on page 165.
Support for large files and large file systems	VxFS supports files larger than two gigabytes and large file systems up to 256 terabytes. Warning: Some applications and utilities might not work on large files. See “largefiles and nolargefiles mount options” on page 167.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
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 <code>getfacl(1)</code> and <code>setfacl(1)</code> manual pages.
Storage Checkpoints	To increase availability, recoverability, and performance, Veritas File System (VxFS) 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 388.
FileSnaps	A FileSnap is a space-optimized copy of a file in the same name space, stored in the same file system. VxFS supports FileSnaps on file systems with disk layout Version 8 or later. See “About FileSnaps” on page 327.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
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 <code>find</code> and <code>cpio</code>) ■ Backing up the entire file system (using <code>fscat</code>) ■ Initiating a full or incremental backup (using <code>vxdump</code>) See “About snapshot file systems” on page 330.
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 693.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
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 SFCFSHA to provide a complete clustering environment. See the <i>Veritas Storage Foundation Cluster File System High Availability Administrator's Guide</i>. To be a cluster mount, a file system must be mounted using the <code>mount -o cluster</code> option. File systems mounted without the <code>-o cluster</code> option are termed local mounts. See the <code>mount_vxfs(1M)</code> manual page. SFCFSHA functionality is separately licensed.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
Support for databases	Databases can be created on files exported as character devices to achieve the same performance as databases created on raw disks. 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. ■ <code>read()</code> and <code>write()</code> 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.
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 <i>Veritas Storage Foundation and High Availability Solutions Solutions Guide</i>.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
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 700.
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 “About reverse path name lookup” on page 708.
Multi-volume file systems	The multi-volume file system (MVFS) 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. MVFS functionality is a separately licensed feature. See “About multi-volume file systems” on page 459.
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 474.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
Thin Reclamation	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. See “About Thin Reclamation of a file system” on page 433.
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 <code>vxtunefs(1M)</code> and <code>fsadm_vxfs(1M)</code> manual pages.
Data deduplication	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. This feature is available only for Solaris SPARC. See “About deduplicating data” on page 560.

Table 2-1 Veritas File System features (*continued*)

Feature	Description
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 571.

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 44.

VxFS provides the following performance enhancements:

- Data synchronous I/O
See [“Data synchronous I/O”](#) on page 295.
- Direct I/O and discovered direct I/O
See [“Direct I/O”](#) on page 294.
See [“Discovered Direct I/O”](#) on page 295.
- Delayed allocation for extending writes
See [“Delayed allocation for extending writes”](#) on page 43.
- Enhanced I/O performance
See [“Enhanced I/O performance”](#) on page 42.
- Caching advisories
See [“Cache advisories”](#) on page 297.
- Enhanced directory features
- Explicit file alignment, extent size, and preallocation controls
- Tunable I/O parameters

- Integration with Veritas Volume Manager (VxVM)
See [“About Veritas Volume Manager”](#) on page 29.
- Support for large directories

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

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

Enhanced I/O performance

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

See [“Enhanced I/O clustering”](#) on page 42.

See [“Veritas Volume Manager integration with Veritas File System for enhanced I/O performance”](#) on page 42.

See [“Application-specific parameters for enhanced I/O performance”](#) on page 43.

Enhanced I/O clustering

I/O clustering is a technique of grouping multiple I/O operations together for improved performance. Veritas File System (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

Veritas File System (VxFS) interfaces with Veritas Volume Manager (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 `vxtunefs` command. You can display the delayed allocation range in the file by using the `fsmap` command.

See the `vxtunefs(1M)` and `fsmap(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](#)

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

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.

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

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 components 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 116.

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.

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 Veritas File System (VxFS).

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

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

How Veritas Volume Manager works with the operating system

Veritas Volume Manager (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. See the <code>vxconfigd(1m)</code> manual page.
<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. See the <code>vxiod(1m)</code> manual page.
<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. See the <code>vxrelocd(1m)</code> manual page.

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 Veritas Volume Manager handles storage management”](#) on page 49.

How Veritas Volume Manager handles storage management

Veritas Volume Manager (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 49.
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 51.

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 3-1](#) shows how a physical disk and device name (*devname*) are illustrated in this document.

Figure 3-1 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.

About disk partitions

[Figure 3-2](#) shows how a physical disk can be divided into one or more slices, also known as partitions.

Figure 3-2 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 3-3 shows how VxVM represents the disks in a disk array as several volumes to the operating system.

Figure 3-3 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 53.
- VM disks
See [“VM disks”](#) on page 54.
- Subdisks
See [“Subdisks”](#) on page 55.
- Plexes
See [“Plexes”](#) on page 56.

- Volumes

See “[Volumes](#)” on page 57.

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 Veritas Volume Manager

Veritas Volume Manager (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 3-4](#) shows the connections between VxVM virtual objects and how they relate to physical disks.

Figure 3-4 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. Disk group names must not contain periods (.).

See [“VM disks”](#) on page 54.

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 588.

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 53.

[Figure 3-5](#) shows a VM disk with a media name of `disk01` that is assigned to the physical disk, *devname*.

Figure 3-5 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 53.

Figure 3-6 shows `disk01-01` is the name of the first subdisk on the VM disk named `disk01`.


Figure 3-6 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 3-7 shows a VM disk with three subdisks, which are assigned from one physical disk.

Figure 3-7 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

Veritas Volume Manager (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 3-8](#) shows an example of a plex with two subdisks.

Figure 3-8 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 layouts.

See [“Volume layouts in Veritas Volume Manager”](#) on page 59.

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 3-9 shows a volume `vol01` with a single plex.

Figure 3-9 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 3-10 shows a mirrored volume `vol06` with two data plexes.

Figure 3-10 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 66.

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

See [“About layered volumes”](#) on page 74.

About the configuration daemon in Veritas Volume Manager

The Veritas Volume Manager (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 `vxdtctl` command is the command-line interface to the `vxconfigd` daemon.

You can use `vxdtctl` 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/rmdp` 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 100.

Volume layouts in Veritas Volume Manager

A Veritas Volume Manager (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 [“About layered volumes”](#) on page 74.

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 61.
- Striping (RAID-0)
See [“Striping \(RAID-0\)”](#) on page 63.
- Mirroring (RAID-1)

See [“Mirroring \(RAID-1\)”](#) on page 66.

- Striping plus mirroring (mirrored-stripe or RAID-0+1)
See [“Striping plus mirroring \(mirrored-stripe or RAID-0+1\)”](#) on page 67.
- 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 68.
- RAID-5 (striping with parity)
See [“RAID-5 \(striping with parity\)”](#) on page 69.

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 3-11](#) 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 3-11 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 3-12](#) shows data spread over two subdisks in a spanned plex.

Figure 3-12 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 3-13 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 3-14 shows a striped plex with three equal sized, single-subdisk columns.

Figure 3-14 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 3-15](#) shows a striped plex with three columns containing subdisks of different sizes.

Figure 3-15 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 145.

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 143.


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 3-16](#) 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 3-16 Mirrored-stripe volume laid out on six disks



See [“Creating a mirrored-stripe volume”](#) on page 145.

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)

Veritas Volume Manager (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 [“About layered volumes”](#) on page 74.

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 3-17](#) 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 3-17 Striped-mirror volume laid out on six disks



See [“Creating a striped-mirror volume”](#) on page 146.

[Figure 3-18](#) 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 3-18 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 3-19](#) shows parity locations in a RAID-5 array configuration.

Figure 3-19 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 71.


Note: Veritas Volume Manager (VxVM) supports RAID-5 for private disk groups, but not for shareable disk groups in a Cluster Volume Manager (CVM) environment. In addition, VxVM does not support the mirroring of RAID-5 volumes that are configured using VxVM 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 3-20 shows the row and column arrangement of a traditional RAID-5 array.

Figure 3-20 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 (VxVM) differs from the traditional structure. Due to the virtual nature of its disks and other objects, VxVM does not use rows.

Figure 3-21 shows how VxVM uses columns consisting of variable length subdisks, where each subdisk represents a specific area of a disk.

Figure 3-21 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 63.

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

See [“Creating a RAID-5 volume”](#) on page 146.

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 3-22 shows a left-symmetric parity layout with five disks (one per column).

Figure 3-22 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 3-23 shows a RAID-5 volume configured across three disks (A, B, and C).

Figure 3-23 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.

About layered volumes

A layered volume is a virtual Veritas Volume Manager (VxVM) 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 3-24 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 146. See the <code>vxassist(1M)</code> manual page.
Creating concatenated-mirrors	See “Creating a concatenated-mirror volume” on page 144. See the <code>vxassist(1M)</code> manual page.
Online Relayout	See “Online relayout” on page 76. 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 85. 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 605.

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 3-25](#) shows how decreasing the number of columns can require disks to be added to a volume.

Figure 3-25 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 3-26](#) shows an example of applying relayout a RAID-5 volume.

Figure 3-26 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 3-27](#) shows an example.

Figure 3-27 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 3-28 shows an example of changing the number of columns.

Figure 3-28 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 3-29 shows an example of changing the column stripe width.

Figure 3-29 Example of increasing the stripe width for the columns in a volume



See [“Performing online relayout”](#) on page 605.

See [“Permitted relayout transformations”](#) on page 606.

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, `bootdkg`), 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 82.

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

See [“SmartSync recovery accelerator”](#) on page 83.

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 544.

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 83.

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 3-30](#) shows how a snapshot volume represents a copy of an original volume at a given point in time.

Figure 3-30 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.

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 3-1](#) compares the features of the various types of snapshots that are supported in VxVM.

Table 3-1 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
Requires less storage space than original volume	No	Yes	No

Table 3-1 Comparison of snapshot features for supported snapshot types
(continued)

Snapshot feature	Full-sized instant (vxsnap)	Space-optimized instant (vxsnap)	Break-off (vxassist or vxsnap)
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 Veritas Volume Manager (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 89.
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 90.

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 588.

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 93.

[Figure 3-31](#) shows an example of a mirrored volume with two plexes on which persistent FastResync is enabled.


Figure 3-31 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 3-32](#) shows how a snapshot plex is set up in the volume, and how a disabled DCO plex is associated with it.


Figure 3-32 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 3-33](#) shows how the creation of the snapshot volume also sets up a DCO object and a DCO volume for the snapshot volume.

Figure 3-33 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 82.

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 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 a 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 that gets allocated during the first write on that node. By default, the size of the DCO volume accommodates 32 DRL maps, an accumulator, and 16 per-volume maps (including a DRL recovery map, a detach map to track detached plexes, and the 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 Veritas Volume Manager (VxVM) that allow several volumes to be represented by a single logical object. All I/O from and to the underlying volumes is directed by way of 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 452.

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 Veritas Volume Manager (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 (`sald`) on the primary SAL server without needing to specify the `sal_username` attribute to `vxassist`.

[Figure 3-34](#), shows how you might choose to set up storage groups within a SAN.

Figure 3-34 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 arrays and on different geographic locations.

The `vxassist` utility in VxVM 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 Veritas Dynamic Multi-Pathing works

This chapter includes the following topics:

- [How DMP works](#)
- [Veritas Volume Manager co-existence with Oracle Automatic Storage Management 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 187.

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 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 4-1 shows how DMP sets up a node for a disk in a supported disk array.


Figure 4-1 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 4-2 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 4-2 Example of multi-pathing for a disk enclosure in a SAN environment



See [“About enclosure-based naming”](#) on page 104.

See [“Changing the disk device naming scheme”](#) on page 264.

See [“Discovering and configuring newly added disk devices”](#) on page 187.

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 193.

About enclosure-based naming

Enclosure-based naming provides an alternative to operating system-based device naming. This allows disk devices to be named for enclosures rather than for the controllers through which they are accessed. 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. 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. Enclosure-based naming allows SF 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.

[Figure 4-3](#) shows a typical SAN environment where host controllers are connected to multiple enclosures through a Fibre Channel switch.

Figure 4-3 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 SF assigns to an enclosure with one that is more meaningful to your configuration.

See [“Renaming an enclosure”](#) on page 242.

[Figure 4-4](#) 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 4-4 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 264.

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 (`error`) that performed error processing, and another (`restore`) 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 `restore` 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 247.

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 107.

See [“I/O throttling”](#) on page 108.

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 the response to I/O failures”](#) on page 243.

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 I/O 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 Subpaths Failover Groups \(SFG\)”](#) on page 246.

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 a subpaths failover group (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 Low Impact Path Probing”](#) on page 245.

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 I/O throttling mechanism”](#) on page 244.

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 I/O policy”](#) on page 234.

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 207.

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 Veritas Volume Manager (VxVM). Clustering is only supported for VxVM.

In a clustered environment where Active/Passive (A/P) type disk arrays are shared by multiple hosts, all nodes in the cluster must access the disk through the same physical storage controller port. Accessing a disk through 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 Array Support Library (ASL) is installed on the system. An Array Policy Module (APM) may also be required.

See [“Discovering disks and dynamically adding disk arrays”](#) on page 189.

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 through the same physical path.

See [“How to administer the Device Discovery Layer”](#) on page 193.

See [“Configuring array policy modules”](#) on page 249.

About enabling or disabling controllers with shared disk groups

Prior to release 5.0, Veritas Volume Manager (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 disks

Automatic Storage Management (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: `cxytdz`

where *partition* is a partition name in the format: `cxytdzsn`

- 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 5. Provisioning new storage](#)
- [Chapter 6. Advanced allocation methods for configuring storage](#)
- [Chapter 7. Creating and mounting VxFS file systems](#)
- [Chapter 8. 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 SF 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”](#) on page 422.

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 Veritas Volume Manager (VxVM) aware of the new LUN size.

```
# vxdisk -g dg1 resize c0t1d0s4
```

See [“Dynamic LUN expansion”](#) on page 257.

- 3 Calculate the new maximum volume size:

```
# vxassist -g dg1 -b maxgrow voll
```

- 4 Grow the volume and the file system to the desired size:

```
# vxresize -b -F vxfs -g dg1 voll 100g
```

Displaying SF information with vxlist

The `vxlist` command is a display command that provides a consolidated view of the SF 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/dsk/clt3d0s1	/
fs	mnt	vxfs	19.84g	9.96g	49%	/dev/vx/dsk/bardg/voll	/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 ”](#) on page 422.

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 120.

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 584.

If you want to assign particular characteristics for a certain volume, you can specify additional attributes on the `vxassist` command line. These can be storage specifications to select certain types of disks for allocation, or other attributes such as the 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 141.
Media types	See “Creating volumes on specific media types” on page 150.
Specific disks, subdisks, plexes locations	See “Creating a volume on specific disks” on page 148.
Ordered allocation	See “Specifying ordered allocation of storage to volumes” on page 150.
Site-based allocation	See “Site-based allocation” on page 153.
Setting the read policy	See “Changing the read policy for mirrored volumes” on page 153.

The `vxassist` utility also provides various constructs to help define and manage volume allocations, with efficiency and flexibility.

See [“Setting default values for vxassist”](#) on page 120.

See [“Using rules to make volume allocation more efficient”](#) on page 121.

See [“Understanding persistent attributes”](#) on page 124.

See [“Customizing disk classes for allocation”](#) on page 127.

See [“Specifying allocation constraints for vxassist operations with the use clause and the require clause”](#) on page 129.

See [“Management of the use and require type of persistent attributes ”](#) on page 137.

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
```

Using rules 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, the volume is 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`.

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 that 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	TUTIL0	PUTIL0
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

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 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.

You can use the `persist` attribute to record allocation attributes on the command line or in a rule file.

See [“Using persistent attributes”](#) on page 125.

You can manage the use and require type of persistent attributes with the intent management operations: `setrule`, `changerule`, `clearrule`, and `listrule`.

See [“Management of the use and require type of persistent attributes ”](#) on page 137.

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	-	-	-

Customizing disk classes for allocation

The `vxassist` command accepts disk classes to indicate storage specifications for allocation. The disk classes are internally-discovered attributes that are automatically associated with the disks. You can specify disk classes to an allocation request with `vxassist` to indicate the type of storage to allocate.

For more information about the disk classes, see the Storage Specifications section of the `vxassist(1M)` manual page.

You can customize the disk classes in the following ways:

- Create a customized alias name.
See [“User-defined alias names for disk classes”](#) on page 127.
- Customize the priority order for the disk classes.
See [“User-defined precedence order for disk classes”](#) on page 127.

You can also create customized, user-defined disk classes.

See [“User-defined disk classes”](#) on page 128.

User-defined alias names for disk classes

For convenience, you can define alias names for existing storage-specification disk classes. Typically, an alias is a shorter or more user-friendly name. You can use the alias instead of its corresponding disk class, to specify `vxassist` allocation constraints. Define the alias names in rule files.

For example, to define “atyp” as an alias for the base disk class “arraytype”, include the following statement in a rule file.

```
class alias atyp=arraytype
```

When the above rule file is used, you can specify the alias “atyp” for allocation. For example, the following constraint specification allocates storage from A/A arrays for the volume creation.

```
# vxassist -g dgname make volname volsize use=atyp:A/A
```

User-defined precedence order for disk classes

The `vxassist` command applies a default priority order for the disk classes that are specified in the mirror confinement (`mirrorconfine`, `wantmirrorconfine`), mirror separation (`mirror`, `wantmirror`), and stripe separation (`stripe`, `wantstripe`) constraints. The higher priority class is honored for allocation when mirroring or striping. If a different priority order is required, you can change the default order for these disk classes.

Note: The “site” class always has the highest precedence, and its order cannot be overridden.

Define the customized precedence order in a rule file. The higher the order number, the higher is the class precedence.

The following shows the default precedence order, for the class names supported with mirror and stripe separation or confinement constraints.

site	order=1000
vendor	order=900
arrayproduct	order=800
array	order=700
arrayport	order=600
hostport	order=400

The acceptable range for the precedence order is between 0 and 1000.

For example, the array class has a higher priority than the hostport class by default. To make the hostport class have a higher priority, assign the hostport class a higher order number. To define the order for the classes, include the following statement in a rule file:

```
class define array order=400
class define hostport order=700
```

When the above rule is used, the following command mirrors across hostport class rather than the array class.

```
# vxassist -g dgname make volname volsize mirror=array,hostport
```

User-defined disk classes

You can define customized disk classes to use in storage specifications for the `vxassist` command. Customized disk classes allow for user-defined device classification and grouping. You can use these disk classes to control allocations. A customized disk class is a user-defined property that is associated with a set of disks. The property is attached as a disk class to the disks that satisfy a particular constraint.

You can use the custom disk classes like other storage-specification disk classes, to specify `vxassist` allocation constraints. Define the custom disk classes in a rule file.

Example

With the following definition in the rule file, the user-defined property “poolname” is associated to the referenced disks. All devices that have the array vendor property defined as HITACHI or IBM, are marked as poolname “finance”. All devices that have the array vendor property defined as DGC or EMC, are marked as poolname “admin”.

```
disk properties vendor:HITACHI {  
    poolname:finance  
}  
disk properties vendor:IBM {  
    poolname:finance  
}  
disk properties vendor:DGC {  
    poolname:admin  
}  
disk properties vendor:EMC {  
    poolname:admin  
}
```

You can now use the user-defined disk class “poolname” for allocation. For example, the following constraint specification allocates disks from the poolname “admin” for the volume creation.

```
# vxassist -g dgname make volname volsize poolname:admin
```

Specifying allocation constraints for vxassist operations with the use clause and the require clause

The `vxassist` command accepts a variety of storage specifications for allocations. The `require` constraint and the `use` constraint are methods to specify detailed storage specifications for allocations. These constraints enable you to select disks from an intersection set or a union set of intended properties. You can specify the set of disks for allocations with more precision than the existing methods `alloc` and `logdisk` clauses. The `use` and `require` constraints can apply to data, log, or both data and log.

The constraints can be one of the following types:

- The `require` constraints

All of the specifications in the constraint must be satisfied, or the allocation fails. A require constraint behaves as an intersection set. For example, allocate disks from a particular array vendor AND with a particular array type.

- The use constraints
At least one of the specifications in the constraint must be satisfied, or the allocation fails. A use constraint behaves as a union set. For example, allocate disks from any of the specified enclosures: `enclrA` or `enclrB`.

For disk group version of 180 or above, the use and require type of constraints are persistent for the volume by default. The default preservation of these clauses enables further allocation operations like grow, without breaking the specified intents.

You can specify multiple storage specifications, separated by commas, in a use or require clause on the `vxassist` command line. You can also specify multiple use or require clauses on the `vxassist` command line.

See [“Interaction of multiple require and use constraints”](#) on page 131.

Use the `vxassist` intent management operations (`setrule`, `changerule`, `clearrule`, `listrule`) to manage persistent require and use constraints.

See [“Management of the use and require type of persistent attributes ”](#) on page 137.

About require constraints

The “require” type of constraints specify that the allocation must select storage that matches all the storage specifications in the constraint. Therefore, the require constraint acts like an intersection set, or a logical AND operation. If any of the specifications cannot be met, the operation fails. The attribute names to specify require constraints are:

- `require`
The constraint applies to both data and log allocation.
- `logrequire`
The constraint applies to log allocations only.
- `datarequire`
The constraint applies to data allocations only.

If any storage-specification is negated with `!`, the allocation excludes the storage that matches that storage specification

Note: If the require type of constraint is provided with the same class but different instances, then these instances are unionized rather than intersected. That is, the allocation selects storage that satisfies any of these storage specifications (similar to use type of constraint).

See [“Interaction of multiple require and use constraints”](#) on page 131.

About use constraints

The “use” type of constraints specify that the allocation must select storage that matches at least one of the storage specifications in the constraint. Therefore, the use constraint acts like a union set, or a logical OR operation. If none of the specifications can be met, the operation fails. The attribute names to specify use constraints are:

- `use`
The constraint applies to both data and log allocation.
- `loguse`
The constraint applies to log allocations only.
- `datause`
The constraint applies to data allocations only.

See [“Interaction of multiple require and use constraints”](#) on page 131.

If the storage specification is negated with `!`, then the allocation excludes the storage that matches that storage specification.

Interaction of multiple require and use constraints

You can specify multiple use or require clauses on the `vxassist` command line. Not all combinations are supported. However, all possible constraint specifications can be achieved with the supported combinations.

The scope for a constraint can be data-specific (`datause` or `datarequire`), log-specific (`loguse` or `logrequire`) or general, which applies to both data and log (`use` or `require`).

Note: Symantec recommends that you do not combine use or require constraints with direct storage-specifications or other clauses like `alloc` or `logdisk`.

The following rules apply when multiple use or require clauses are specified:

- Multiple use constraints of the same scope are unionized, so that at least one of the storage specifications is satisfied. That is, multiple `use` clauses; multiple `datause` clauses; or multiple `loguse` clauses.
- Multiple require constraints of the same scope are intersected, so that all the storage specifications are satisfied. That is, multiple `require` clauses; multiple `datarequire` clauses; or multiple `logrequire` clauses.

- Require and use constraints of the same scope are mutually intersected. That is, `require` clauses and `use` clauses; `datarequire` clauses and `datause` clauses; or `logrequire` clauses and `loguse` clauses. At least one of the use storage specifications must be satisfied and all of the require storage specifications are satisfied. For example, if a `datause` clause and a `datarequire` clause are used together, the allocation for the data must meet at least one of the `datause` specifications and all of the `datarequire` specifications.
- Data-specific constraints and log-specific constraints can be used together. They are applied independently to data and logs respectively. That is, `datause` clause with `loguse` clause or `logrequire` clause; `datarequire` clause with `loguse` clause or `logrequire` clause. For example, a `datarequire` clause can be used to control data allocation along with a `logrequire` clause to control log allocation.
- The `vxassist` command does not support a mix of general scope constraints with data-specific or log-specific constraints. For example, a `require` clause cannot be used along with the `logrequire` clause or a `datarequire` clause. However, all possible constraint specifications can be achieved with the supported combinations.

[Table 6-1](#) summarizes these rules for the interaction of each type of constraint if multiple constraints are specified.

Table 6-1 Combinations of require and use constraints

Scope	Mutually unionized	Mutually intersected	Applied independently
Data	<code>datause</code> - <code>datause</code>	<code>datarequire</code> - <code>datause</code> <code>datarequire</code> - <code>datarequire</code>	<code>datause</code> - <code>loguse</code> <code>datause</code> - <code>logrequire</code> <code>datarequire</code> - <code>loguse</code> <code>datarequire</code> - <code>logrequire</code>
Log	<code>loguse</code> - <code>loguse</code>	<code>logrequire</code> - <code>loguse</code> <code>logrequire</code> - <code>logrequire</code>	<code>loguse</code> - <code>datause</code> <code>loguse</code> - <code>datarequire</code> <code>logrequire</code> - <code>datause</code> <code>logrequire</code> - <code>datarequire</code>
General - log and data	<code>use</code> - <code>use</code>	<code>use</code> - <code>require</code> <code>require</code> - <code>require</code>	N/A

Examples of use and require constraints

The following examples show use and require constraints for storage allocation.

Example 1 - require constraint

This example shows the require constraint in a disk group that has disks from two arrays: `emc_clariion0` and `ams_wms0`. Both arrays are connected through the same HBA hostportid (06-08-02), but the arrays have different arraytype (A/A and A/A-A respectively).

The following output shows the disk group information:

```
# vxprint -g testdg
TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg testdg        testdg          -        -        -        -        -        -

dm ams_wms0_359  ams_wms0_359 -        2027264 -        -        -        -
dm ams_wms0_360  ams_wms0_360 -        2027264 -        -        -        -
dm ams_wms0_361  ams_wms0_361 -        2027264 -        -        -        -
dm ams_wms0_362  ams_wms0_362 -        2027264 -        -        -        -
dm emc_clariion0_0 emc_clariion0_0 -    4120320 -        -        -        -
dm emc_clariion0_1 emc_clariion0_1 -    4120320 -        -        -        -
dm emc_clariion0_2 emc_clariion0_2 -    4120320 -        -        -        -
dm emc_clariion0_3 emc_clariion0_3 -    4120320 -        -        -        -
```

To allocate both the data and the log on the disks that are attached to the particular HBA and that have the array type A/A:

```
# vxassist -g testdg make v1 1G logtype=dc0 dcoverversion=20 \
require=hostportid:06-08-02,arraytype:A/A
```

The following output shows the results of the above command. The command allocated disk space for the data and the log on `emc_clariion0` array disks, which satisfy all the storage specifications in the `require` constraint:

```
# vxprint -g testdg
TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg testdg        testdg          -        -        -        -        -        -

dm ams_wms0_359  ams_wms0_359 -        2027264 -        -        -        -
dm ams_wms0_360  ams_wms0_360 -        2027264 -        -        -        -
dm ams_wms0_361  ams_wms0_361 -        2027264 -        -        -        -
dm ams_wms0_362  ams_wms0_362 -        2027264 -        -        -        -
dm emc_clariion0_0 emc_clariion0_0 -    4120320 -        -        -        -
dm emc_clariion0_1 emc_clariion0_1 -    4120320 -        -        -        -
dm emc_clariion0_2 emc_clariion0_2 -    4120320 -        -        -        -
```

```

dm emc_clariion0_3 emc_clariion0_3 - 4120320 - - - -
v v1 fsgen ENABLED 2097152 - ACTIVE - -
pl v1-01 v1 ENABLED 2097152 - ACTIVE - -
sd emc_clariion0_0-01 v1-01 ENABLED 2097152 0 - - -
dc v1_dco v1 - - - - - - -
v v1_dcl gen ENABLED 67840 - ACTIVE - -
pl v1_dcl-01 v1_dcl ENABLED 67840 - ACTIVE - -
sd emc_clariion0_0-02 v1_dcl-01 ENABLED 67840 0 - - -

```

Example 2 - use constraint

This example shows the use constraint in a disk group that has disks from three arrays: `ams_wms0`, `emc_clariion0`, and `hitachi_vsp0`.

The following output shows the disk group information:

```

# vxprint -g testdg
TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg testdg        testdg          -        -        -        -        -        -

dm ams_wms0_359 ams_wms0_359 -        2027264 -        -        -        -        -
dm ams_wms0_360 ams_wms0_360 -        2027264 -        -        -        -        -
dm ams_wms0_361 ams_wms0_361 -        2027264 -        -        -        -        -
dm ams_wms0_362 ams_wms0_362 -        2027264 -        -        -        -        -
dm emc_clariion0_0 emc_clariion0_0 - 4120320 -        -        -        -        -
dm hitachi_vsp0_3 hitachi_vsp0_3 - 4120320 -        -        -        -        -

```

To allocate both the data and the log on the disks that belong to the array `ams_wms0` or the array `emc_clariion0`:

```

# vxassist -g testdg make v1 3G logtype=dco dcoverion=20 \
use=array:ams_wms0,array:emc_clariion0

```

The following output shows the results of the above command. The command allocated disk space for the data and the log on disks that satisfy the arrays specified in the `use` constraint.

```

# vxprint -g testdg
TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg testdg        testdg          -        -        -        -        -        -

dm ams_wms0_359 ams_wms0_359 -        2027264 -        -        -        -        -
dm ams_wms0_360 ams_wms0_360 -        2027264 -        -        -        -        -
dm ams_wms0_361 ams_wms0_361 -        2027264 -        -        -        -        -
dm ams_wms0_362 ams_wms0_362 -        2027264 -        -        -        -        -

```

```

dm emc_clariion0_0 emc_clariion0_0 - 4120320 - - -
dm hitachi_vsp0_3 hitachi_vsp0_3 - 4120320 - - -

v v1 fsgen ENABLED 6291456 - ACTIVE - -
pl v1-01 v1 ENABLED 6291456 - ACTIVE - -
sd ams_wms0_359-01 v1-01 ENABLED 2027264 0 - - -
sd ams_wms0_360-01 v1-01 ENABLED 143872 2027264 - - -
sd emc_clariion0_0-01 v1-01 ENABLED 4120320 2171136 - - -
dc v1_dco v1 - - - - - -
v v1_dcl gen ENABLED 67840 - ACTIVE - -
pl v1_dcl-01 v1_dcl ENABLED 67840 - ACTIVE - -
sd ams_wms0_360-02 v1_dcl-01 ENABLED 67840 0 - - -

```

Example 3: datause and logrequire combination

This example shows the combination of a `datause` constraint and a `logrequire` constraint. The disk group has disks from three arrays: `ams_wms0`, `emc_clariion0`, and `hitachi_vsp0`, which have different array types.

The following output shows the disk group information:

```

# vxprint -g testdg
TY NAME ASSOC KSTATE LENGTH PLOFFS STATE TUTILO PUTILO
dg testdg testdg - - - - -
dm ams_wms0_359 ams_wms0_359 - 2027264 - - -
dm ams_wms0_360 ams_wms0_360 - 2027264 - - -
dm ams_wms0_361 ams_wms0_361 - 2027264 - - -
dm ams_wms0_362 ams_wms0_362 - 2027264 - - -
dm emc_clariion0_0 emc_clariion0_0 - 4120320 - - -
dm emc_clariion0_1 emc_clariion0_1 - 4120320 - - -
dm emc_clariion0_2 emc_clariion0_2 - 4120320 - - -
dm emc_clariion0_3 emc_clariion0_3 - 4120320 - - -
dm hitachi_vsp0_3 hitachi_vsp0_3 - 4120320 - - -

```

To allocate data on disks from `ams_wms0` or `emc_clariion0` array, and to allocate log on disks from arraytype `A/A-A`:

```

# vxassist -g testdg make v1 1G logtype=dco dconversion=20 \
datause=array:ams_wms0,array:emc_clariion0 logrequire=arraytype:A/A-A

```

The following output shows the results of the above command. The command allocated disk space for the data and the log independently. The data space is allocated on `emc_clariion0` disks that satisfy the `datause` constraint. The log space is allocated on `ams_wms0` disks that are `A/A-A` arraytype and that satisfy the `logrequire` constraint:

```
# vxprint -g testdg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	testdg	testdg	-	-	-	-	-	-
dm	ams_wms0_359	ams_wms0_359	-	2027264	-	-	-	-
dm	ams_wms0_360	ams_wms0_360	-	2027264	-	-	-	-
dm	ams_wms0_361	ams_wms0_361	-	2027264	-	-	-	-
dm	ams_wms0_362	ams_wms0_362	-	2027264	-	-	-	-
dm	emc_clariion0_0	emc_clariion0_0	-	4120320	-	-	-	-
dm	emc_clariion0_1	emc_clariion0_1	-	4120320	-	-	-	-
dm	emc_clariion0_2	emc_clariion0_2	-	4120320	-	-	-	-
dm	emc_clariion0_3	emc_clariion0_3	-	4120320	-	-	-	-
dm	hitachi_vsp0_3	hitachi_vsp0_3	-	4120320	-	-	-	-
v	v1	fsgen	ENABLED	2097152	-	ACTIVE	-	-
pl	v1-01	v1	ENABLED	2097152	-	ACTIVE	-	-
sd	emc_clariion0_0-01	v1-01	ENABLED	2097152	0	-	-	-
dc	v1_dco	v1	-	-	-	-	-	-
v	v1_dcl	gen	ENABLED	67840	-	ACTIVE	-	-
pl	v1_dcl-01	v1_dcl	ENABLED	67840	-	ACTIVE	-	-
sd	ams_wms0_359-01	v1_dcl-01	ENABLED	67840	0	-	-	-

Example 4 - use and require combination

This example shows the combination of a `use` constraint and a `require` constraint.

The disk group has disks from three arrays: `ams_wms0`, `emc_clariion0`, and `hitachi_vsp0`. Only the disks from `ams_wms0` array are multi-pathed.

The following output shows the disk group information:

```
# vxprint -g testdg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	testdg	testdg	-	-	-	-	-	-
dm	ams_wms0_359	ams_wms0_359	-	2027264	-	-	-	-
dm	ams_wms0_360	ams_wms0_360	-	2027264	-	-	-	-
dm	ams_wms0_361	ams_wms0_361	-	2027264	-	-	-	-
dm	ams_wms0_362	ams_wms0_362	-	2027264	-	-	-	-
dm	emc_clariion0_0	emc_clariion0_0	-	4120320	-	-	-	-
dm	emc_clariion0_1	emc_clariion0_1	-	4120320	-	-	-	-
dm	emc_clariion0_2	emc_clariion0_2	-	4120320	-	-	-	-
dm	emc_clariion0_3	emc_clariion0_3	-	4120320	-	-	-	-
dm	hitachi_vsp0_3	hitachi_vsp0_3	-	4120320	-	-	-	-

To allocate data and log space on disks from `emc_clariion0` or `ams_wms0` array, and disks that are multi-pathed:


```
# vxassist -g testdg make v1 1G logtype=dc0 dcoverversion=20 \
use=array:emc_clariion0,array:ams_wms0 require=multipathed:yes
```

The following output shows the results of the allocation. The data and log space is on `ams_wms0` disks, which satisfy the `use` as well as the `require` constraints:

```
# vxprint -g testdg
```

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE	TUTILO	PUTILO
dg	testdg	testdg	-	-	-	-	-	-
dm	ams_wms0_359	ams_wms0_359	-	2027264	-	-	-	-
dm	ams_wms0_360	ams_wms0_360	-	2027264	-	-	-	-
dm	ams_wms0_361	ams_wms0_361	-	2027264	-	-	-	-
dm	ams_wms0_362	ams_wms0_362	-	2027264	-	-	-	-
dm	emc_clariion0_0	emc_clariion0_0	-	4120320	-	-	-	-
dm	emc_clariion0_1	emc_clariion0_1	-	4120320	-	-	-	-
dm	emc_clariion0_2	emc_clariion0_2	-	4120320	-	-	-	-
dm	emc_clariion0_3	emc_clariion0_3	-	4120320	-	-	-	-
dm	hitachi_vsp0_3	hitachi_vsp0_3	-	4120320	-	-	-	-
v	v1	fsgen	ENABLED	2097152	-	ACTIVE	-	-
pl	v1-01	v1	ENABLED	2097152	-	ACTIVE	-	-
sd	ams_wms0_359-01	v1-01	ENABLED	2027264	0	-	-	-
sd	ams_wms0_360-01	v1-01	ENABLED	69888	2027264	-	-	-
dc	v1_dco	v1	-	-	-	-	-	-
v	v1_dcl	gen	ENABLED	67840	-	ACTIVE	-	-
pl	v1_dcl-01	v1_dcl	ENABLED	67840	-	ACTIVE	-	-
sd	ams_wms0_360-02	v1_dcl-01	ENABLED	67840	0	-	-	-

Management of the use and require type of persistent attributes

Persistent attributes are the saved vlume intents that should be honored for subsequent allocation operations for that volume. The intent management operations enable you to manage the use and require type of persistent intents for volumes. These operations allow you to independently manage the intents after the volume creation. When you change the persistent intents for a volume, the changed intents are not checked for validity or enforced for the current allocation of the volume.

You can set, change, clear, or list the persistent intents for the volume with the following `vxassist` operations:

- `setrule`
Replaces any existing saved intents with the specified intents for the specified volume.
- `changerule`

Appends the specified intents to the existing saved intents for the specified volume.

- `clearrule`

Removes any existing saved intents for the specified volume.

- `listrule`

Lists any saved intents for the specified volume. If no volume name is specified, the command shows the intents for all of the volumes.

The intent management operations only apply to the use or require type of persistent constraints. The other type of persistent constraints are managed with the `persist` attribute.

See [“Using persistent attributes”](#) on page 125.

To display the intents that are currently associated to a volume

- ◆ To display the intents that are currently associated to a volume, use the following command:

```
# vxassist [options] listrule [volume]
```

For example, to display the existing saved intents for the volume `v1`:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    require=array:ams_wms0
}
```

To replace the intents that are currently associated to a volume

- 1 Display the intents that are currently associated to the volume:

```
# vxassist [options] listrule [volume]
```

In this example, the volume v1 has an existing saved intent that requires the array to be `ams_wms0`. For example, to display the existing saved intents for the volume v1:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    require=array:ams_wms0
}
```

- 2 Specify the new intent with the following command:

```
# vxassist [options] setrule volume attributes...
```

For example, to replace the array with the `ds4100-0` array, specify the new intent with the following command:

```
# vxassist -g testdg setrule v1 require=array:ds4100-0
```

- 3 Verify the new intent with the display command.

For example, the following command shows that the intent has changed:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    require=array:ds4100-0
}
```

To add to the intents that are currently associated to a volume

- 1 Display the intents that are currently associated to the volume:

```
# vxassist [options] listrule [volume]
```

In this example, the volume v1 has an existing saved intent that requires the array to be ds4100-0. For example, to display the existing saved intents for the volume v1:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    use=array:ds4100-0
}
```

- 2 Add the new intent with the following command:

```
# vxassist [options] changerule volume attributes...
```

For example, to add the `ams_wms0` array in the use constraint, specify the new intent with the following command:

```
# vxassist -g testdg changerule v1 use=array:ams_wms0
```

- 3 Verify the new intent with the display command.

For example, the following command shows that the intent has changed:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    use=array:ds4100-0,array:ams_wms0
}
```

To clear the intents that are currently associated to a volume

- 1 Display the intents that are currently associated to the volume:

```
# vxassist [options] listrule [volume]
```

For example, to display the existing saved intents for the volume v1:

```
# vxassist -g testdg listrule v1
volume rule v1 {
    require=multipathed:yes
    use=array:emc_clariion0,array:ams_wms0
}
```

- 2 Clear the existing intents with the following command:

```
# vxassist [options] clearrule volume
```

For example, to clear the intents for the volume v1:

```
# vxassist -g testdg clearrule v1
```

- 3 Verify that the volume has no saved intents.

For example, the following command shows that the volume v1 has no saved intents:

```
# vxassist -g testdg listrule v1
volume rule v1 {}
```

Creating volumes of a specific layout

Veritas Volume Manager (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:

- mirrored volumes
See [“Creating a mirrored volume”](#) on page 143.
- striped volumes
See [“Creating a striped volume”](#) on page 145.
- RAID-5 volumes
See [“Creating a RAID-5 volume”](#) on page 146.

Types of volume layouts

Veritas Volume Manager (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 61.
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 63.
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 66.
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 69.

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 67.
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 “About layered volumes” on page 74. 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 68. ■ 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 611.

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 63.

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 611.

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 69.

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 150.

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 to select a particular type of disks for a certain volume, you can provide the storage specifications to `vxassist` for storage allocation.

For more information, see the Storage Specifications section of the `vxassist(1M)` manual page.

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 150.

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 97.

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 6-1](#) shows an example of using ordered allocation to create a mirrored-stripe volume.

Figure 6-1 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 6-2](#) shows an example of using ordered allocation to create a striped-mirror volume.

Figure 6-2 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 6-3 shows an example of using concatenated disk space to create a mirrored-stripe volume.

Figure 6-3 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 6-4 shows an example of using storage allocation to create a mirrored-stripe volume across controllers.

Figure 6-4 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.

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.
<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:

```
/opt/VRTS/bin/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 Veritas File System disk layouts”](#) on page 750.

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:

```
# /opt/VRTS/bin/mkfs /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/dsk/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

You can mount a VxFS file system by using the `mount` command. When you enter the `mount` command, the generic `mount` command parses the arguments and the `-F FSType` option executes the `mount` command specific to that file system type. The `mount` command first searches the `/etc/fs/FSType` directory, then the `/usr/lib/fs/FSType` directory. If the `-F` option is not supplied, the command searches the file `/etc/vfstab` for a file system and an `FSType` matching the special file or mount point provided. If no file system type is specified, `mount` uses the default file system.

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.

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 and nolargefiles mount options](#)

- [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.

To mount a file system

- ◆ Use the `mount` command to mount a file system:

```
mount [-F vxfs] [generic_options] [-r] [-o specific_options] \
special mount_point
```

<code>vxfs</code>	File system type.
<code>generic_options</code>	Options common to most other file system types.
<code>specific_options</code>	Options specific to VxFS.
<code>-o ckpt=ckpt_name</code>	Mounts a Storage Checkpoint.
<code>-o cluster</code>	Mounts a file system in shared mode. Available only with the VxFS cluster file system feature.
<code>special</code>	A VxFS block special device.

<i>mount_point</i>	Directory on which to mount the file system.
<i>-r</i>	Mounts the file system as read-only.

log mount option

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. 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 which 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. 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.

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 `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. `mwdisable` 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 and nolargefiles mount options

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 167.

See [“Mounting a file system with large files”](#) on page 168.

See [“Managing a file system with large files”](#) on page 168.

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 323.

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/c1t3d0s1 /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 Veritas File System disk layouts”](#) on page 750.

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
```

`vxfs`

The file system type.

<i>newsize</i>	The size to which the file system will increase. The default units is sectors, but you can specify <i>k</i> or <i>K</i> for kilobytes, <i>m</i> or <i>M</i> for megabytes, or <i>g</i> or <i>G</i> for gigabytes.
<i>mount_point</i>	The file system's mount point.
<i>-r rawdev</i>	Specifies the path name of the raw device if there is no entry in <i>/etc/vfstab</i> and <i>fsadm</i> 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
```

vxfs The file system type.

<i>newsize</i>	The size to which the file system will shrink. The default units is sectors, but you can specify <i>k</i> or <i>K</i> for kilobytes, <i>m</i> or <i>M</i> for megabytes, or <i>g</i> or <i>G</i> for gigabytes.
<i>mount_point</i>	The file system's mount point.
<i>-r rawdev</i>	Specifies the path name of the raw device if there is no entry in <i>/etc/vfstab</i> and <i>fsadm</i> 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
```

<i>vxfs</i>	The file system type.
<i>-d</i>	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.
<i>-D</i>	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/voll` 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/voll`:

```
# fstyp -v /dev/vx/dsk/fsvol/voll
```

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 irstart 0
bstart 34 femap 1051 fimap 0 fiextop 0 fistart 0 fbstart
```

Monitoring free space

In general, Veritas File System (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.

VxFS supports reclamation of free storage on a Thin Storage LUN.

See [“About Thin Reclamation of a file system”](#) on page 433.

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
```

Extent attributes

This chapter includes the following topics:

- [About extent attributes](#)
- [Commands related to extent attributes](#)

About extent attributes

Veritas File System (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.

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 `cp_vxfs(1)`, `cpio_vxfs(1)`, `mv_vxfs(1)`, `getext(1)`, `setext(1)`, `vxdump(1M)`, and `vxrestore(1M)` manual pages.

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 180.

See [“Fixed extent size”](#) on page 180.

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 181.

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 [“About Veritas File System I/O”](#) on page 293.

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 182.
- Whether allocations are contiguous
See [“Extent attribute contiguity”](#) on page 182.
- Whether the file can be written beyond its reservation
See [“Write operations beyond extent attribute reservation”](#) on page 182.

- Whether an unused reservation is released when the file is closed
See [“Extent attribute reservation trimming”](#) on page 182.
- Whether the reservation is a persistent attribute of the file
See [“Extent attribute reservation persistence”](#) on page 183.
- 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 183.

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 293.

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 9. Administering Dynamic Multi-Pathing](#)
- [Chapter 10. Dynamic Reconfiguration of devices](#)
- [Chapter 11. Managing devices](#)
- [Chapter 12. 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 DMP database 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 SF's configuration daemon was last started. These disks are then dynamically configured and recognized by SF.

`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 Dynamic Multi-Pathing (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 SF 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 SF 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 `c1t1d0s2` and `c2t2d0s2`:

```
# vxdisk scandisks device=c1t1d0s2,c2t2d0s2
```

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=c1t1d0s2,c2t2d0s2
```

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 200.

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 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 200.

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 version. Be sure to install the VRTSaslapm package that supports the installed version of Storage Foundation.

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 DDL device discovery and includes the new disks in the VxVM device list.

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 SF 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 SF 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. The TPD coexistence feature of SF permits coexistence without requiring any change in a third-party multi-pathing driver.

See [“Changing device naming for TPD-controlled enclosures”](#) on page 267.

See [“Displaying information about TPD-controlled devices”](#) on page 220.

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 204.
- 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 9-1](#) 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 9-1 Scenarios for using DMP with PowerPath

PowerPath	DMP	Array configuration mode
Installed.	The <code>libvxpp</code> ASL handles EMC Symmetrix arrays and DGC CLARiiON claiming internally. PowerPath handles failover.	EMC Symmetrix - Any DGC CLARiiON - 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
Not installed; the array is DGC CLARiiON (CXn00).	DMP handles multi-pathing. The ASL name is <code>libvxCLARiiON</code> .	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 SF 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.
- List all the Host Bus Adapters including iSCSI
- List the ports configured on a Host Bus Adapter
- List the targets configured from a Host Bus Adapter
- List the devices configured from a Host Bus Adapter
- Get or set the iSCSI operational parameters
- List the types of arrays that are supported.
- Add support for an array to DDL.
- Remove support for an array from DDL.
- List information about excluded disk arrays.
- List disks that are supported in the `DISKS` (JBOD) category.
- Add disks from different vendors to the `DISKS` category.
- Remove disks from the `DISKS` category.
- Add disks as foreign devices.

The following 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 9-2 Parameters for iSCSI devices

Parameter	Default value	Minimum value	Maximum value
DataPDUInOrder	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
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
-----	-----	-----	-----	-----
DataPDUInOrder	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 100.

Warning: This procedure ensures that Dynamic Multi-Pathing (DMP) is set up correctly on an array that is not supported by Veritas Volume Manager. Otherwise, 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 `vxdctl enable` command to bring the array under VxVM control.

```
# vxdctl enable
```

See [“Enabling discovery of new disk arrays”](#) on page 191.

- 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	CabinetNum	Policy
		(Cmd/PageCode/off/len)	(Cmd/PageCode/off/len)	
=====				
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, use either the `vxdisk scandisks` or the `vxctl enable` command to discover the devices as simple disks. These disks then behave in the same way as autoconfigured disks.

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 191.

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 SF. 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, array ports, or DMP nodes”](#) on page 240.

See [“Enabling I/O for paths, controllers, array ports, or DMP nodes”](#) on page 241.

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.

About displaying DMP database 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.

See [“Administering DMP using vxddmpadm”](#) on page 211.

Displaying the paths to a disk

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:

```
# vxdisk list devicename
```

For example, to view multi-pathing information for `c2t0d0s2`, use the following command:

```
# vxdisk list c2t0d0s2
```

The output from the `vxdisk list` command displays the multi-pathing information, as shown in the following example:

```
Device      c2t0d0
devicetag   c2t0d0
type        sliced
hostid      sys1
.
.
.
Multipathing information:
numpaths:   2
c2t0d0s2    state=enabled   type=primary
c1t0d0s2    state=disabled  type=secondary
```

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
```

Typical output from the `vxddmpadm getsubpaths` command is as follows:

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 213.
- Display consolidated information about the DMP nodes
See [“Displaying consolidated information about the DMP nodes”](#) on page 214.
- Display the members of a LUN group.
See [“Displaying the members of a LUN group”](#) on page 215.
- 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 215.
- Display information about the HBA controllers on the host.
See [“Displaying information about controllers”](#) on page 218.
- Display information about enclosures.
See [“Displaying information about enclosures”](#) on page 219.
- Display information about array ports that are connected to the storage processors of enclosures.

- See [“Displaying information about array ports”](#) on page 220.
- Display information about devices that are controlled by third-party multi-pathing drivers.
See [“Displaying information about TPD-controlled devices”](#) on page 220.
 - Display extended devices attributes.
See [“Displaying extended device attributes”](#) on page 222.
 - Suppress or include devices from VxVM control.
See [“Suppressing or including devices from VxVM control”](#) on page 224.
 - Gather I/O statistics for a DMP node, enclosure, path or controller.
See [“Gathering and displaying I/O statistics”](#) on page 225.
 - Configure the attributes of the paths to an enclosure.
See [“Setting the attributes of the paths to an enclosure”](#) on page 230.
 - Display the redundancy level of a device or enclosure
See [“Displaying the redundancy level of a device or enclosure”](#) on page 232.
 - Specify the minimum number of active paths
See [“Specifying the minimum number of active paths”](#) on page 233.
 - Display or set the I/O policy that is used for the paths to an enclosure.
See [“Specifying the I/O policy”](#) on page 234.
 - Enable or disable I/O for a path, HBA controller or array port on the system.
See [“Disabling I/O for paths, controllers, array ports, or DMP nodes”](#) on page 240.
 - Rename an enclosure.
See [“Renaming an enclosure”](#) on page 242.
 - Configure how DMP responds to I/O request failures.
See [“Configuring the response to I/O failures”](#) on page 243.
 - Configure the I/O throttling mechanism.
See [“Configuring the I/O throttling mechanism”](#) on page 244.
 - Control the operation of the DMP path restoration thread.
See [“Configuring DMP path restoration policies”](#) on page 247.
 - Configure array policy modules
See [“Configuring array policy modules”](#) on page 249.
 - Get or set the values of various tunables used by DMP.
See [“DMP tunable parameters”](#) on page 713.

The following sections cover these tasks in detail along with sample output.

See the `vxdkmpadm(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 `vxddmpadm 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:

```
# vxddmpadm list dmpnode all
```

Use the `enclosure` attribute with `list dmpnode` to obtain a list of all DMP nodes for the specified enclosure.

```
# vxddmpadm list dmpnode enclosure=enclosure name
```

For example, the following command displays the consolidated information for all of the DMP nodes in the `enc0` enclosure.

```
# vxddmpadm 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 `vxddmpadm disable` command is listed as `disabled(m)`.

```
# vxddmpadm list dmpnode dmpnodename=dmpnodename
```

For example, the following command displays the consolidated information for the DMP node `emc_clariion0_158`.

```
# vxddmpadm 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 aportID aportWWN attr
path          = c0t5006016141E03B33dls2 enabled(a) primary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 A5 50:06:01:61:41:e0:3b:33 -
path          = c0t5006016041E03B33dls2 enabled(a) primary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 A4 50:06:01:60:41:e0:3b:33 -
path          = c0t5006016841E03B33dls2 enabled secondary FC c0
/pci@1e,600000/SUNW,emlxs@3/fp@0,0 B4 50:06:01:68:41:e0:3b:33 -
path          = c1t5006016141E03B33dls2 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          = c1t5006016841E03B33dls2 enabled secondary FC c1
/pci@1e,600000/SUNW,emlxs@3,1/fp@0,0 B4 50:06:01:68:41:e0:3b:33 -
path          = c1t5006016041E03B33dls2 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:

```
# vxddmpadm getlungroup dmpnodename=c11t0d10s2
```

The above command displays output such as the following:

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/rmp` 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 I/O policy”](#) on page 234.

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:

```
# vxddmpadm getsubpaths ctlr=c2
```

NAME	STATE [-]	PATH-TYPE [-]	CTLR-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:

```

# vxddmpadm getsubpaths enclosure=enclosure portid=portid
# vxddmpadm getsubpaths pwwn=pwwn

```

For example, to list subpaths through an array port through the enclosure and the array port ID:

```
# vxddmpadm getsubpaths enclosure=HDS9500-ALUA0 portid=1A
```

```

NAME          STATE[A]    PATH-TYPE[M]  DMPNODENAME  ENCLR-NAME    CTLR  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:

```
# vxddmpadm getsubpaths pwwn=20:00:00:E0:8B:06:5F:19
```

```

NAME          STATE[A]    PATH-TYPE[M]  DMPNODENAME  ENCLR-NAME    CTLR  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.

```
# vxddmpadm 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:

```
# vxdkmpadm getsubpaths enclosure=Disk ctrl=c1
```

NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-NAME	CTRL	ATTRS
c1t65d0s2	ENABLED (A)	-	Disk_1	Disk	c1	-
c1t66d0s2	ENABLED (A)	-	Disk_2	Disk	c1	-

By default, the output of the `vxdkmpadm 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:

```
# vxdkmpadm -s {path | dmpnode | enclosure | ctrl} getsubpaths \
[all | ctrl=ctrl_name | dmpnodename=dmp_device_name | \
enclosure=enclr_name [ctrl=ctrl_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:

```
# vxdkmpadm listctrl all
```

CTRL-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 `vxdkmpadm 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	qllogic	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 information about TPD-controlled devices

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:

```
# vxdmppadm getsubpaths tpdnodename=TPD_node_name
# vxdmppadm gettpdnode nodename=TPD_path_name
```

See [“Changing device naming for TPD-controlled enclosures”](#) on page 267.

For example, consider the following disks in an EMC Symmetrix array controlled by PowerPath, which are known to DMP:

```
# vxdisk 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`:

```
# vxdmppadm getsubpaths tpdnodename=emcpower10s2
```

NAME	TPDNODENAME	PATH-TYPE[-]	DMP-NODENAME	ENCLR-TYPE	ENCLR-NAME
=====					
c7t0d10s2	emcpower10s2-		emcpower10s2	EMC	EMC0
c6t0d10s2	emcpower10s2-		emcpower10s2	EMC	EMC0

Conversely, the next command displays information about the PowerPath node that corresponds to the path, `c7t0d10s2`, discovered by DMP:

```
# vxdmppadm 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                DDL_DEVICE_ATTR  VID
tagmastore-usp0_0a7a  std fc RAID_5    HITACHI
tagmastore-usp0_0a7b  std fc RAID_5    HITACHI
tagmastore-usp0_0a78  std fc RAID_5    HITACHI
tagmastore-usp0_0a79  std fc RAID_5    HITACHI
tagmastore-usp0_065a  hdprclm fc          HITACHI
tagmastore-usp0_065b  hdprclm fc          HITACHI
tagmastore-usp0_065c  hdprclm fc          HITACHI
tagmastore-usp0_065d  hdprclm fc          HITACHI

```

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  cl0t0d3         std fc RAID_5
tagmastore-usp0_0a78  auto  -    -      online  cl0t0d0         std fc RAID_5
tagmastore-usp0_0655  auto  -    -      online  cl3t2d7         hdprclm fc

```

```
tagmastore-usb0_0656 auto - - online c13t3d0 hdprclm fc
tagmastore-usb0_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 `vxdkmpadm 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 `vxdkmpadm 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.

```
# vxdkmpadm exclude { all | product=VID:PID |
ctrlr=[\!]ctrlrname | dmpnodename=diskname [ path=[\!]pathname] }

# vxdkmpadm include { all | product=VID:PID |
ctrlr=[\!]ctrlrname | dmpnodename=diskname [ path=[\!]pathname] }
```

where:

`all` – all devices

`product=VID:PID` – all devices with the specified VID:PID

`ctrlr=ctrlrname` – 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 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 | ctrlr=ctrlr-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 I/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]
```

For example:

```
# 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
```

	OPERATIONS		BLOCKS		AVG TIME (ms)	
CTRLNAME	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:

```
# vxddpdm -u m iostat show groupby=arrayport \
enclosure=HDS9500-ALUA0 portid=1A
```

	OPERATIONS		BYTES		AVG TIME (ms)	
PORTNAME	READS	WRITES	READS	WRITES	READS	WRITES
1A	743	1538	11m	24m	17.13	8.61

To group by enclosure:

```
# vxddpdm [-u unit] iostat show groupby=enclosure [ all \
| enclosure=enclr ]
```

For example:

```
# vxddpdm -u h iostat show groupby=enclosure enclosure=EMC_CLARiiON0
```

	OPERATIONS		BLOCKS		AVG TIME (ms)	
ENCLNAME	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:

```
# vxddpdm [-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:

```
# vxddpdm -z iostat show dmpnodename=c2t16d4s2
```

	OPERATIONS		BLOCKS		AVG TIME (ms)	
PATHNAME	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.

```
# vxddpdm -u us iostat show pathname=c2t17d4s2
```

	OPERATIONS		BLOCKS		AVG TIME (microsec)	
PATHNAME	READS	WRITES	READS	WRITES	READS	WRITES
c2t17d4s2	20	126	4	29	9500.00	19413.79

Displaying 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:

```
# vxddmpadm -q iostat show dmpnodename=c5t2d1s2
```

	QUEUED I/Os		Pending I/Os
DMPNODENAME	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:

```
# 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 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.

The attributes set for the paths are persistent and are stored in the `/etc/vx/dmppolicy.info` file.

You can set the following attributes:

active	Changes a standby (failover) path to an active path. The following example specifies an active path for an array: <pre># vxddmpadm setattr path c2t10d0s2 pathtype=active</pre>
nomanual	Restores the original primary or secondary attributes of a path. This example restores the path to a JBOD disk: <pre># vxddmpadm setattr path c3t10d0s2 pathtype=nomanual</pre>
nopreferred	Restores the normal priority of a path. The following example restores the default priority to a path: <pre># vxddmpadm setattr path c1t20d0s2 \ pathtype=nopreferred</pre>
preferred [priority=N]	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. See “Specifying the I/O policy” on page 234. This example first sets the I/O policy to <code>priority</code> for an Active/Active disk array, and then specifies a preferred path with an assigned priority of 2: <pre># vxddmpadm setattr enclosure enc0 \ iopolicy=priority # vxddmpadm setattr path c1t20d0s2 \ pathtype=preferred priority=2</pre>
primary	Defines a path as being the primary path for a JBOD disk array. The following example specifies a primary path for a JBOD disk array: <pre># vxddmpadm setattr path c3t10d0s2 \ pathtype=primary</pre>
secondary	Defines a path as being the secondary path for a JBOD disk array. The following example specifies a secondary path for a JBOD disk array: <pre># vxddmpadm setattr path c4t10d0s2 \ pathtype=secondary</pre>

standby 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. The next example specifies a standby path for an A/P-C disk array:

```
# vxddmpadm setattr path c2t10d0s2 \
    pathtype=standby
```

Displaying the redundancy level of a device or enclosure

Use the `vxddmpadm getdmpnode` 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 getdmpnode enclosure=enc1_name redundancy=value
```

For example, to list the devices with fewer than 3 enabled paths, use the following command:

```
# vxddmpadm getdmpnode 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

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 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 `partitionsizes` 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 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: 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:

<code>adaptive</code>	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 <code>enc1</code>: <pre># vxddmpadm setattr enclosure enc1 \ iopolicy=adaptive</pre>
<code>adaptiveminq</code>	Similar to the <code>adaptive</code> policy, except that I/O is scheduled according to the length of the I/O queue on each path. The path with the shortest queue is assigned the highest priority.

```
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). Specifying a partition size of 0 is equivalent to the default partition size of 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 713.

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:

```
# vxdkmpadm 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 230.

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 or an ALUA 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 or an ALUA 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 or ALUA.

For example, the following command sets the `balanced` I/O policy with a partition size of 4096 blocks (2MB) on the enclosure `enc0`, and allows scheduling of I/O requests on the secondary paths:

```
# 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	ATTR_NAME	DEFAULT	CURRENT
=====			
HDS9500-ALUA0	use_all_paths	no	yes

The `use_all_paths` attribute only applies to A/A-A arrays and ALUA 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 this sample output from the `vxdisk list` command, the device `c3t2d15s2` 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 &
```

By running the `vxdkmpadm iostat` command to display the DMP statistics for the device, it can be seen that all I/O is being directed to one path, `c5t4d15s2`:

```
# 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

The vxdkmpadm command is used to display the I/O policy for the enclosure that contains the device:

```
# vxdkmpadm getattr enclosure ENC0 iopolicy
```

ENCLR_NAME	DEFAULT	CURRENT
=====		
ENC0	MinimumQ	Single-Active

This 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, the policy is set to round-robin as shown here:

```
# vxdkmpadm setattr enclosure ENC0 iopolicy=round-robin
# vxdkmpadm getattr enclosure ENC0 iopolicy
```

ENCLR_NAME	DEFAULT	CURRENT
=====		
ENC0	MinimumQ	Round-Robin

The DMP statistics are now reset:

```
# vxdkmpadm iostat reset
```

With the workload still running, the effect of changing the I/O policy to balance the load across the primary paths can now be seen.

```
# vxdkmpadm iostat show dmpnodename=c3t2d15s2 interval=5 count=2
.
.
.
```

cpu usage = 14403us per cpu memory = 32768b

PATHNAME	OPERATIONS		KBYTES		AVG TIME (ms)	
	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

The enclosure can be returned to the single active I/O policy by entering the following command:

```
# vxdkmpadm setattr enclosure ENC0 iopolicy=singleactive
```

Disabling I/O for paths, controllers, array ports, or DMP nodes

Disabling I/O through a path, HBA controller, array port, or DMP node prevents DMP from issuing I/O requests through the specified path, or the paths that are connected to the specified controller, array port, or DMP node. 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:


```
# vxddmpadm [-c|-f] disable enclosure=enclr_name portid=array_port_ID
# vxddmpadm [-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:

```
# vxddmpadm disable enclosure=HDS9500V0 portid=1A
# vxddmpadm 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:

```
# vxddmpadm [-c|-f] disable ctrlr=ctrlr_name enclosure=enclr_name \
portid=array_port_ID
```

To disable I/O for a particular DMP node, specify the DMP node name.

```
# vxddmpadm [-c|-f] disable dmpnodename=dmpnode
```

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, array ports, or DMP nodes

Enabling a controller allows a previously disabled path, HBA controller, array port, or DMP node to accept I/O again. This operation succeeds only if the path, controller, array port, or DMP node 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:

```
# vxdkmpadm 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:

```
# vxdkmpadm enable ctlr=ctlr_name1[,ctlr_name2,ctlr_nameN]
```

To enable I/O for the paths connected to an array port, use one of the following commands:

```
# vxdkmpadm enable enclosure=enclr_name portid=array_port_ID
```

```
# vxdkmpadm 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:

```
# vxdkmpadm enable enclosure=HDS9500V0 portid=1A
```

```
# vxdkmpadm 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:

```
# vxdkmpadm enable ctlr=ctlr_name enclosure=enclr_name \
  portid=array_port_ID
```

To enable I/O for a particular DMP node, specify the DMP node name.

```
# vxdkmpadm enable dmpnodename=dmpnode
```

Renaming an enclosure

The `vxdkmpadm setattr` command can be used to assign a meaningful name to an existing enclosure, for example:

```
# vxdkmpadm 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:

```
# vxdkmpadm listenclosure all
```

ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS

```
=====
other0      OTHER      OTHER_DISKS      CONNECTED
jbod0       X1         X1_DISKS        CONNECTED
GRP1        ACME        60020f20000001a90000  CONNECTED
```

Configuring the 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 246.

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 I/O throttling mechanism”](#) on page 244.

Note: The response to I/O failure settings is persistent across reboots of the system.

Configuring the 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 246.

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 the response to I/O failures”](#) on page 243.

Note: The I/O throttling settings are persistent across reboots of the system.

Configuring Low Impact Path Probing

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 Subpaths Failover Groups (SFG)

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 9-3](#) summarizes the possible recovery option settings for retrying I/O after an error.

Table 9-3 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 9-4](#) summarizes the possible recovery option settings for throttling I/O.

Table 9-4 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 249.

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 713.

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](#)
- [Replacing a host bus adapter online](#)
- [Upgrading the array controller firmware online](#)

About online Dynamic Reconfiguration

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.

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 253.
- Replacing a host bus adapter (HBA) online
See [“Replacing a host bus adapter online”](#) on page 260.
- Updating the array controller firmware, also known as a nondisruptive upgrade
See [“Upgrading the array controller firmware online”](#) on page 260.

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.

The operations are as follows:

- Dynamic LUN removal from an existing target ID
See [“Removing LUNs dynamically from an existing target ID”](#) on page 253.
- Dynamic new LUN addition to a new target ID
See [“Adding new LUNs dynamically to a new target ID”](#) on page 255.
- Replacing a LUN on an existing target ID
See [“Replacing LUNs dynamically from an existing target ID”](#) on page 256.
- Dynamic LUN expansion
See [“Dynamic LUN expansion”](#) on page 257.
- Changing the LUN characteristics
See [“Changing the characteristics of a LUN from the array side”](#) on page 259.

Removing LUNs dynamically from an existing target ID

Veritas Dynamic Multi-Pathing (DMP) provides a Dynamic Reconfiguration tool to simplify the removal of LUNs from an existing target ID. Each LUN is unmapped from the host. DMP issues an operating system device scan and cleans up the operating system device tree.

Warning: Do not run any device discovery operations outside of the Dynamic Reconfiguration tool until the device operation is completed.

To remove LUNs dynamically from an existing target ID

- 1 Remove the device from use by any volume manager.

If the device is in use by Veritas Volume Manager (VxVM), perform the following steps:

- If the device is part of a disk group, move the disk out of the disk group.

```
# vxdg -g dgname rmdisk daname
```

See [“Removing a disk from a disk group”](#) on page 626.

- Remove the disk from the `vxdisk` list.
In a cluster, perform this step from all of the nodes.

```
# vxdisk rm da-name
```

For example:

```
# vxdisk rm eva4k6k0_0
```

For LUNs that are in use by ZFS, export or destroy the zpool.

- 2 Start the `vxdiskadm` utility:

```
# vxdiskadm
```

- 3 Select the **Dynamic Reconfiguration operations** option from the `vxdiskadm` menu.
- 4 Select the **Remove LUNs** option.
- 5 Type **list** or press **Return** to display a list of LUNs that are available for removal. A LUN is available for removal if it is not in a disk group and the state is online, nolabel, online invalid, or online thinrlm.

The following shows an example output:

```
Select disk devices to remove: [<pattern-list>,all,list]: list
LUN(s) available for removal:
eva4k6k0_0
eva4k6k0_1
eva4k6k0_2
eva4k6k0_3
eva4k6k0_4
emc0_017e
```

- 6 Enter the name of a LUN, a comma-separated list of LUNs, or a regular expression to specify the LUNs to remove.

For example, enter `emc0_017e`.

- 7 At the prompt, confirm the LUN selection.

DMP removes the LUN from VxVM usage.

- 8 At the following prompt, remove the LUN from the array.

```
Enclosure=emc0 AVID=017E
Device=emc0_017e Serial=830017E000
-----
PATH=c15t0d6 ctrlr=c15 port=7e-a [50:01:43:80:12:08:3c:26]
PATH=c17t0d6 ctrlr=c17 port=7e-a [50:01:43:80:12:08:3a:76]
-----
Please remove LUNs with Above details from array and
press 'y' to continue removal (default:y):
```

- 9 Return to the Dynamic Reconfiguration tool and select **y** to continue the removal process.

DMP completes the removal of the device from VxVM usage. Output similar to the following displays:

```
Luns Removed
-----
emc0_017e
```

DMP updates the operating system device tree and the VxVM device tree.

- 10 Select **exit** to exit the Dynamic Reconfiguration tool.

Adding new LUNs dynamically to a new target ID

The Dynamic Reconfiguration tool simplifies the addition of new LUNs to an existing target ID. One or more new LUNs are mapped to the host by way of multiple HBA ports. An operating system device scan is issued for the LUNs to be recognized and added to DMP control.

Warning: Do not run any device discovery operations outside of the Dynamic Reconfiguration tool until the device operation is completed.

To add new LUNs dynamically to a new target ID

- 1 Start the `vxdiskadm` utility:

```
# vxdiskadm
```

- 2 Select the **Dynamic Reconfiguration operations** option from the `vxdiskadm` menu.

- 3 Select the **Add LUNs** option.

The tool issues a device discovery.

- 4 When the prompt displays, add the LUNs from the array.

- 5 Select **y** to continue to add the LUNs to DMP.

The operation issues a device scan. The newly-discovered devices are now visible.

```
Luns Added
-----
Enclosure=emc0 AVID=017E
Device=emc0_017e Serial=830017E000
PATH=c15t0d6 ctrlr=c15 port=7e-a [50:01:43:80:12:08:3c:26]
PATH=c17t0d6 ctrlr=c17 port=7e-a [50:01:43:80:12:08:3a:76]
```

- 6 Select a label to add to the disk.
- 7 Select **exit** to exit the Dynamic Reconfiguration tool.

Replacing LUNs dynamically from an existing target ID

The Veritas Dynamic Multi-Pathing (DMP) provides a Dynamic Reconfiguration tool to simplify the replacement of new LUNs from an existing target ID. Each LUN is unmapped from the host. DMP issues an operating system device scan and cleans up the operating system device tree.

Warning: Do not run any device discovery operations outside of the Dynamic Reconfiguration tool until the device operation is completed.

To replace LUNs dynamically from an existing target ID

- 1 Remove the device from use by any volume manager.

If the device is in use by Veritas Volume Manager (VxVM), perform the following steps:

- If the device is part of a disk group, move the disk out of the disk group.

```
# vxdg -g dgname rmdisk daname
```

See [“Removing a disk from a disk group”](#) on page 626.

- Remove the disk from the `vxdisk` list.

In a cluster, perform this step from all of the nodes.

```
# vxdisk rm da-name
```


For example:

```
# vxdisk rm eva4k6k0_0
```

For LUNs that are in use by ZFS, export or destroy the zpool.

- 2 Start the `vxdiskadm` utility:

```
# vxdiskadm
```

- 3 Select the **Dynamic Reconfiguration operations** option from the `vxdiskadm` menu.

- 4 Select the **Replace LUNs** option.

The output displays a list of LUNs that are available for replacement. A LUN is available for replacement if it is not in a disk group and the state is online, nolabel, online invalid, or online thinrcm.

- 5 Select one or more LUNs to replace.

- 6 At the prompt, confirm the LUN selection.

- 7 Remove the LUN from the array.

- 8 Return to the Dynamic Reconfiguration tool and select `y` to continue the removal.

After the removal completes successfully, the Dynamic Reconfiguration tool prompts you to add a LUN.

- 9 When the prompt displays, add the LUNs from the array.

- 10 Select `y` to continue to add the LUNs to DMP.

The operation issues a device scan. The newly-discovered devices are now visible.

DMP updates the operating system device tree.

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.

Changing the characteristics of a LUN from the array side

Some arrays provide a way to change the properties of LUNs. For example, the EMC Symmetrix array allows write-protected (Read-only), and read-write enabled LUNs. Before changing the properties of a LUN, you must remove the device from Veritas Volume Manager (VxVM) control.

To change the properties of a LUN

- 1 Remove the device from use by any volume manager.

If the device is in use by Veritas Volume Manager (VxVM), perform the following steps:

- If the device is part of a disk group, move the disk out of the disk group.

```
# vxdg -g dgname rmdisk daname
```

See [“Removing a disk from a disk group”](#) on page 626.

- Remove the disk from the `vxdisk` list.

In a cluster, perform this step from all of the nodes.

```
# vxdisk rm da-name
```

For example:

```
# vxdisk rm eva4k6k0_0
```

For LUNs that are in use by ZFS, export or destroy the zpool.

- 2 Change the device characteristics.
- 3 Use SF to perform a device scan. In a cluster, perform this command on all the nodes.

```
# vxdisk scandisks
```

- 4 Add the device back to the disk group.

```
# vxdg -g dname adddisk daname
```

Replacing a host bus adapter online

Veritas Dynamic Multi-Pathing (DMP) provides a Dynamic Reconfiguration tool to simplify the removal of host bus adapters from an existing system.

To replace a host bus adapter online

- 1 Start the `vxdiskadm` utility:

```
# vxdiskadm
```

- 2 Select the **Dynamic Reconfiguration operations** option from the `vxdiskadm` menu.
- 3 Select the **Replace HBAs** option.
The output displays a list of HBAs that are available to DMP.
- 4 Select one or more HBAs to replace.
- 5 At the prompt, confirm the HBA selection.
- 6 Replace the host bus adapter.
- 7 Return to the Dynamic Reconfiguration tool and select `y` to continue the replacement process.

DMP updates the operating system device tree.

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. SF 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 SF 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 SF 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 267.

Displaying the disk-naming scheme

SF disk naming can be operating-system based naming or enclosure-based naming.

The following command displays whether the SF 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
NAMING_SCHEME      PERSISTENCE LOWERCASE USE_AVID
=====
Enclosure Based   Yes           Yes           Yes
```

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 SF 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 SF 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 265.

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`.

```
# vxdisk 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
```

```
# vxdisk 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 269.

See [“Removing the error state for simple or nopriv disks in non-boot disk groups”](#) on page 270.

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#t#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

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 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 SF 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 SF 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 273.

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 205.

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 273.

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 675.

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 273.

Removing disks

You must disable a disk group before you can remove the last disk in that group.

See [“Disabling a disk group”](#) on page 655.

As an alternative to disabling the disk group, you can destroy the disk group.

See [“Destroying a disk group”](#) on page 655.

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 662.

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](#)
- [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 288.
- Logging of DMP events for troubleshooting (DMP event)
See [“DMP event logging”](#) on page 290.
- Automated device discovery (OS event)
See [“Dynamic Multi-Pathing \(DMP\) automated device discovery”](#) on page 289.
- 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 290.

See “[Starting and stopping the Dynamic Multi-Pathing \(DMP\) event source daemon](#)” on page 291.

Fabric Monitoring and proactive error detection

DMP takes a proactive role in detecting errors on paths. The DMP event source daemon `vxesd` uses the Storage Networking Industry Association (SNIA) HBA API library to receive SAN fabric events from the HBA. DMP checks devices that are suspect based on the information from the SAN events, even if there is no active I/O. New I/O is directed to healthy paths while DMP verifies the suspect devices.

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

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. On Solaris, 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
```

Optimizing I/O performance

- [Chapter 13. Veritas File System I/O](#)
- [Chapter 14. 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 294.

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 169.
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 and High Availability Solutions products improve the overall performance of database environments in a variety of ways.

Table 13-1 Storage Foundation and High Availability Solutions database accelerators

SFHA Solutions database accelerator	Supported databases	Use cases and considerations
Oracle Disk Manager (ODM)	Oracle	■ To improve Oracle performance and manage system bandwidth through an improved Application Programming Interface (API) that contains advanced kernel support for file I/O. ■ To use Oracle Resilvering and turn off Veritas Volume Manager Dirty Region Logging (DRL) to increase performance, use ODM. ■ To reduce the time required to restore consistency, freeing more I/O bandwidth for business-critical applications, use SmartSync recovery accelerator.
Cached Oracle Disk Manager (Cached ODM)	Oracle	To enable selected I/O to use caching to improve ODM I/O performance, use Cached ODM.
Quick I/O (QIO)	Oracle DB2 Sybase	To achieve raw device performance for databases run on VxFS file systems, use Veritas Quick I/O.

Table 13-1 Storage Foundation and High Availability Solutions database accelerators *(continued)*

SFHA Solutions database accelerator	Supported databases	Use cases and considerations
Cached Quick I/O (Cached QIO)	Oracle DB2 Sybase	To further enhance database performance by leveraging large system memory to selectively buffer the frequently accessed data, use Cached QIO.
Concurrent I/O	DB2 Sybase	Concurrent I/O (CIO) is optimized for DB2 and Sybase environments To achieve improved performance for databases run on VxFS file systems without restrictions on increasing file size, use Veritas Concurrent I/O.

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 and High Availability Solutions database accelerators enable you to manage performance for your database with more precision.

For details about using ODM, Cached ODM, QI/O, and Cached QIO for Oracle, see *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*.

For details about using QIO, Cached QIO, and Concurrent I/O for DB2, see *Veritas Storage Foundation: Storage and Availability Management for DB2 Databases*.

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

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`
- `vxsnap addmir`
- `vxsnap reattach`

- vxsd mv

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`.

Using Point-in-time copies

- [Chapter 15. Understanding point-in-time copy methods](#)
- [Chapter 16. Administering volume snapshots](#)
- [Chapter 17. Administering Storage Checkpoints](#)
- [Chapter 18. Administering FileSnaps](#)
- [Chapter 19. 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 15-1](#) illustrates the steps that are needed to set up the processing solution on the primary host.

Figure 15-1


Using snapshots and FastResync to implement point-in-time copy solutions on a primary host



Note: The Disk Group Split/Join functionality is not used. As all processing takes place in the same disk group, synchronization of the contents of the snapshots from the original volumes is not usually required unless you want to prevent disk contention. Snapshot creation and updating are practically instantaneous.

Figure 15-2 shows the suggested arrangement for implementing solutions where the primary host is used and disk contention is to be avoided.

Figure 15-2 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 15-3 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 15-3 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 15-3](#) should be connected to different controllers.

[Figure 15-4](#) shows an example of how you might achieve such connectivity using Fibre Channel technology with 4 Fibre Channel controllers in the primary host.

Figure 15-4 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 15-5 shows how off-host processing might be implemented in a cluster by configuring one of the cluster nodes as the OHP node.

Figure 15-5 Example implementation of an off-host point-in-time copy solution using a cluster node





Figure 15-6 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 15-6 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 15-7](#) illustrates the steps that are needed to set up the processing solution on the primary host.

Figure 15-7 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 a 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 by the FlashSnap 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.
- 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 319.

- File level snapshots.
The FileSnap feature provides snapshots at the level of individual files.

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 15-1

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
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 15-8](#) shows the instant space-optimized snapshot model.

Figure 15-8 Space-optimized instant snapshot creation and usage in a backup cycle



See [“Creating and managing space-optimized instant snapshots”](#) on page 349.

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 15-9 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 15-9 Primary fileset and its Storage Checkpoint



In Figure 15-10, 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 15-9.

Figure 15-10 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 322.

Copy-on-write

In [Figure 15-11](#), 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 15-11 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.

Note: If an auto-mounted Storage Checkpoint is in use by an NFS mount, removing the Storage Checkpoint might succeed even without the forced (`-f`) option.

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.

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 394.

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 15-12](#), 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 15-12 Updates to a nodata clone

See [“Showing the difference between a data and a nodata Storage Checkpoint”](#) on page 394.

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 402.

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.

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

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 297.

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 315.

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`, `snapwait`, 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 16-1 shows the traditional third-mirror break-off volume snapshot model that is supported by the `vxassist` command.

Figure 16-1 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 16-2 shows the full-sized instant volume snapshot model.

Figure 16-2 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 352.

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 347.

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 357.


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 347.

Cascaded snapshots

[Figure 16-3](#) shows a snapshot hierarchy, known as a snapshot cascade, that can improve write performance for some applications.

Figure 16-3 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 363.

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 16-4 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 16-3](#) on page 337.

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 16-5 shows the sequence of steps that would be required to restore a database.

Figure 16-5 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 16-6 shows the effect of applying the `vxsnap dis` command to snapshots with and without dependent snapshots.

Figure 16-6 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 16-7 shows the operation of the `vxsnap split` command.

Figure 16-7 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 16-8 shows an alternative where the snapshot overwrites the original volume.

Figure 16-8 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 366.

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 perform 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 360.

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 360.

See [“Creating and managing space-optimized instant snapshots”](#) on page 349.

See [“Creating and managing full-sized instant snapshots”](#) on page 352.

See [“Creating and managing third-mirror break-off snapshots”](#) on page 354.

See [“Creating and managing linked break-off snapshot volumes”](#) on page 357.

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 add 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 345.

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 347.

Creating a shared cache object

If you need several space-optimized instant snapshots for the volumes in a disk group, you can create a single shared cache object in the disk group rather than a separate cache object for each snapshot.

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 374.

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 148.

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
```

Upgrading the instant snap Data Change Objects (DCOs) and DCO volumes for a VxVM volume

Instant snap DCOs, formerly known as version 20 DCOs, support the creation of instant snapshots for VxVM volumes. Upgrade the instant snap DCOS and DCO volumes to ensure compatibility with the latest version of VxVM. The upgrade operation can be performed while the volumes are online.

The upgrade operation does not support upgrade from version 0 DCOs.

To upgrade the instant snap DCOs for all volumes in the disk group

- 1 Make sure that the disk group is at least version 180. To upgrade the disk group:

```
# vxdg upgrade diskgroup
```

- 2 Use the following command to upgrade the instant snap DCOs for all volumes in the disk group:

```
# vxsnap -g diskgroup upgradeall
```

Where:*diskgroup* is the disk group that contains the volumes to be upgraded.

For additional options to the upgradeall operation, see the `vxsnap(1M)` manual page.

To upgrade the instant snap DCOs for specified volumes

- 1 Make sure that the disk group is at least version 180. To upgrade the disk group:

```
# vxdg upgrade diskgroup
```

- 2 To upgrade the DCOs, specify one or more volumes or volume sets to the following command:

```
# vxsnap [-g diskgroup] upgrade  
[volume1|volset1] [volume2|volset2...]
```

Where:*diskgroup* is the disk group that contains the volumes to be upgraded.

For additional options to the upgrade operation, see the `vxsnap(1M)` manual page.

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 345.

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 364.
 - 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 366.
 - Destroy the snapshot.
See [“Removing an instant snapshot”](#) on page 367.

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 347.

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 370.

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 382.

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 364.

- 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 364.
- 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 366.
- 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 366.
- 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 367.

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 364.
- 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 364.
- 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 366.
- 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 366.
- 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 367.

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 347.

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 364.
 - Reattach the snapshot volume with the original volume.
See [“Reattaching a linked break-off snapshot volume”](#) on page 365.
 - 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 366.
 - 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 367.

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 347.

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 362.

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 370.

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 370.

See [“Removing an instant snapshot”](#) on page 367.

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 374.

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 370.

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 16-1](#) shows the commands that are provided for controlling the synchronization manually.

Table 16-1 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 16-1](#) 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 357.

See [“Reattaching an instant full-sized or plex break-off snapshot”](#) on page 364.

See [“Reattaching a linked break-off snapshot volume”](#) on page 365.

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 383.

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 593.

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 148.

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 378.

- 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 379.

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 380.
- 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 381.
- 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 383.

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 316.

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 383.

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, upgrade 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 384.

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 593.

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 322.

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 320.

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 390.

See [“Removing a Storage Checkpoint”](#) on page 391.

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 324.

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, using the `qiomkfile` 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_fbmap(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 18-1 provides a list of Veritas File System (VxFS) commands that enable you to administer FileSnaps.

Table 18-1

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>fsmap</code>	The <code>fsmap</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>fsmap(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 18-1 (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)    Space_Saved(KB)
/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
```

Comparison of the logical size output of the `fsadm -S shared`, `du`, and `df` commands

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)	Space_Saved (KB)
/mnt	52428800	49073632	83600	83610	10

```
# 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 19-1 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 19-1 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 388.

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 storage with Storage Foundation

- [Chapter 20. Understanding storage optimization solutions in Storage Foundation](#)
- [Chapter 21. Migrating data from thick storage to thin storage](#)
- [Chapter 22. Maintaining Thin Storage with Thin Reclamation](#)

Understanding storage optimization solutions in Storage Foundation

This chapter includes the following topics:

- [About thin provisioning](#)
- [About thin optimization solutions in Storage Foundation](#)
- [About SmartMove](#)
- [About the Thin Reclamation feature](#)
- [Determining when to reclaim space on a thin reclamation LUN](#)
- [How automatic reclamation works](#)

About thin provisioning

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 provides several features to support thin provisioning and thin reclamation, and to optimize storage use on thin provisioned arrays.

See [“About SmartMove”](#) on page 423.

About thin optimization solutions in Storage Foundation

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 work together with the array functionality to solve these challenges.

[Table 20-1](#) lists the Storage Foundation features and benefits relating to thin storage.

Table 20-1 Thin storage solutions in Storage Foundation

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 423. Improves performance for copy operations. Enables migration from thick LUNs to thin provisioned LUNs. See “Migrating to thin provisioning” on page 427.

Table 20-1 Thin storage solutions in Storage Foundation (continued)

Feature	Description	Benefits
Thin disk discovery	Storage Foundation 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 424.

About SmartMove

Storage Foundation 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 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 supports reclamation of the unused storage on thin-reclamation capable arrays. Storage Foundation 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 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 434.
- 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 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 442.

How automatic reclamation works

On thin-reclamable arrays, storage that is no longer in use needs to be reclaimed by the array. Storage Foundation 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 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 442.

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 438.

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 427.

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 614.

- 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](#)
- [About the reclamation log file](#)
- [Monitoring Thin Reclamation using the `vxtask` command](#)
- [Configuring automatic reclamation](#)

Reclamation of storage on thin reclamation arrays

Veritas Storage Foundation supports reclamation of the unused storage on thin-reclamation capable arrays and LUNs. Storage Foundation 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 433.

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 438.
- 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 437.

About Thin Reclamation of a disk, a disk group, or an enclosure

Storage Foundation 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 440.

About 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. 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)` manual page.

Thin Reclamation is only supported on file systems mounted on a VxVM volume.

Thin Reclamation is not supported for file systems mounted on RAID5 volumes.

Veritas File System also supports reclamation of a portion of the file system using the `vxfs_ts_reclaim()` API.

See the `vxfs_ts_reclaim(3)` manual page and 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 automatically discovers thin devices that have been recognized on the host as `thin` or `thinrlm`. DMP uses the Veritas array support libraries (ASLs) to recognize vendor-specific thin attributes and claim devices accordingly as `thin` or `thinrlm`.

Thin devices that are classified as `thin` are capable of thin provisioning. Veritas Thin Reclamation only works on devices with the `thinrlm` attribute set. Before performing thin reclamation, determine whether the system recognizes the LUN as a `thinrlm` host.

To identify devices on a host that are known to have the `thin` or `thinrlm` 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 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 424.

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 `thinreclm` 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 `thinreclm` 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 (thinreclm) LUNs.
GROUP	The disk group that contains the disk.

TYPE The type of thin devices – thin provisioning capable (thin) or thin reclamation capable (thinrclm). The `vxdisk -o thin list` 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 `thinrclm` 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	thinrclm
emc0_428b	16384.00m	3200.00m	22.00m	mydg	thinrclm
emc0_4287	16384.00m	6233.00m	617.00m	mydg	thinrclm
emc0_4288	16384.00m	1584.00m	1417.00m	mydg	thinrclm
emc0_4289	16384.00m	2844.00m	1187.00m	mydg	thinrclm
xiv0_030f	16384.00m	2839.00m	1223.00m	xivdg	thinrclm
xiv0_0307	16384.00m	666.00m	146.00m	xivdg	thinrclm
xiv0_0308	16384.00m	667.00m	147.00m	xivdg	thinrclm
xiv0_0309	16384.00m	3.00m	-	-	thinrclm
xiv0_0310	16384.00m	30.00m	-	-	thinrclm

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	thinrclm
emc0_428b	16384	6335	624	mydg	thinrclm
emc0_4287	16384	6335	617	mydg	thinrclm
emc0_4288	16384	1584	617	mydg	thinrclm
emc0_4289	16384	2844	1187	mydg	thinrclm

Reclaiming space on a file system

Table 22-1 lists the `fsadm` command options that administer thin reclamation.

Table 22-1 `fsadm` options for administering thin reclamation

Option	Description
<code>-o aggressive -A</code>	Initiates Thin Storage aggressive reclamation. Aggressive reclamation is not supported on SSD devices.
<code>-o analyse analyze</code>	Initiates the analyze reclaim option.
<code>-o auto</code>	Initiates the auto reclaim option.
<code>-o ssd</code>	Initiates the TRIM command on an underlying SSD trim-capable device.
<code>-o thin</code>	Initiates thin reclamation on the underlying Thin Reclaim-capable device.
<code>-P</code>	Performs multi-threaded Thin Storage reclamation. By default, the <code>fsadm</code> command performs single-threaded Thin Storage reclamation. To use multi-threaded Thin Storage Reclamation, the array must support multiple concurrent reclaim operations.
<code>-R</code>	Performs reclamation of free storage to the Thin Storage LUN on a VxFS file system .

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

To perform aggressive space reclamation

- 1 Ensure you mounted the VxFS file system.

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

If you must mount the VxFS file system, see the `mount_vxfs(1M)` manual page.

- 2 Perform aggressive reclamation of free storage to the Thin Storage LUN on the VxFS file system that is mounted at `/mnt1`:

```
# /opt/VRTS/bin/fsadm -R -o aggressive /mnt1
```

To perform space reclamation

- 1 Ensure you mounted the VxFS file system.

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

If you must mount the VxFS file system, see the `mount_vxfs(1M)` manual page.

- 2 Perform space reclamation on the VxFS file system that is mounted at `/mnt1`:

```
# /opt/VRTS/bin/fsadm -F vxfs -R /mnt1
```

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.

Storage Foundation logs the statistics for reclamation events in the `/etc/vx/log/reclaim_log` file.

See [“About the reclamation log file”](#) on page 439.

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
```

You can turn off thin reclamation for a specific device with the following command:

```
# vxdisk set reclaim=off disk
```

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

About the reclamation log file

Storage Foundation logs the statistics for reclamation events in the `/etc/vx/log/reclaim_log` file. [Table 22-2](#) describes the fields in the reclamation log file.

For Veritas Volume Replicator (VVR), reclamation logging only happens for the local node.

Table 22-2 The reclamation log file fields

LOG fields	Description
START_TIME	The start time of the reclamation task.
DURATION	The time taken to complete the reclamation task.
DISKGROUP	The disk group name associated with the subdisk. For TYPE=GAP, the disk group value may be NULL value.
VOLUME	The volume associated with the subdisk. If a volume is not associated with the subdisk, the value is NULL.
DISK	The disk associated with the subdisk.
SUBDISK	The subdisk name for which the reclamation operation is performed.
OFFSET	The starting offset of the subdisk.
LEN	The total length of the subdisk.
PA_BEFORE	The physical allocation before the reclamation task.
PA_AFTER	The physical allocation after the reclamation task.
TYPE	The type for the reclamation operation. The value is one of the following: ■ GAP: reclaim the gap between the subdisks ■ SD: reclaim the subdisk ■ FULL: reclaim the full LUN on disk with no DG present ■ VXFS: reclaim a mounted VxFS file system.
STATUS	Whether the reclamation operation succeeded or not. In case of failure, the STATUS also displays the error code. When an object such as a volume or plex is removed, the status is logged as "Pending."

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 23. Understanding storage tiering with SmartTier](#)
- [Chapter 24. Creating and administering volume sets](#)
- [Chapter 25. Multi-volume file systems](#)
- [Chapter 26. Administering SmartTier](#)
- [Chapter 27. Administering hot-relocation](#)
- [Chapter 28. Deduplicating data on Solaris SPARC](#)
- [Chapter 29. Compressing files](#)

Understanding storage tiering with SmartTier

This chapter includes the following topics:

- [About SmartTier](#)
- [How the SmartTier policy works with the shared extents](#)
- [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:

- 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 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

Note: SmartTier is the expanded and renamed feature previously known as Dynamic Storage Tiering (DST).

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.

SmartTier leverages two key technologies included with Veritas Storage Foundation: 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.

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 459.

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.

See [“Converting a single volume file system to a multi-volume file system”](#) on page 463.

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.

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 onlines 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 445.

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 455.
- 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 `voll1`, in the disk group `mydg`, you would use the following command:

```
# vxvset -g mydg -t vxfs make myvset voll1
```

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 457.

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 `mnknd` 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 474.
- 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 461.

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/dg1/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/dg1/myvset /mnt1
# fsvoladm -H list /mnt1
```

devic	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 7
```

- 4 If the disk layout version is 7, upgrade to Version 8:

```
# 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/dg1/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/rdsk/dg1/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/dg1/myvset /mnt1

# mount -F vxfs /dev/vx/dsk/dg1/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 `fsmap` and `fsvmap` commands. The `fsmap`

`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 474.

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](#)
- [Sub-file relocation](#)

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 459.

Note: Some of the commands have changed or been removed between the 4.1 release and the current 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 477.

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>

About compressing files with SmartTier

You can use the SmartTier feature to compress and uncompress files automatically based on the rules defined in a placement policy. SmartTier performs the allocation for compressed or uncompressed extents of the selected files directly from the tier that is specified in the policy. The selected files get compressed or uncompressed while relocating to the specified tier of storage.

You can perform in-place compressing of an entire tier, which compresses all of the uncompressed extents of all of the files on the tier. This operation is useful if a write or append was performed on a file on this tier, which results in the file having some uncompressed extents.

SmartTier uses `gzip` as the default compression algorithm, and 1 MB is the default block size for compression. These default values are not configurable through an XML policy file.

SmartTier can compress and uncompress files as specified by a placement policy in the following ways:

- Compress while relocating files from one tier to another in a multi-volume file system
- Uncompress while relocating files from one tier to another in a multi-volume file system
- Compress in-place in a multi-volume file system
- Uncompress in-place in multi-volume file system
- Compress in-place in single volume file system
- Uncompress in-place in single volume file system
- Compress an entire tier in multi-volume file system
- Uncompress an entire tier in multi-volume file system

See [“About compressing files”](#) on page 571.

Supported SmartTier document type definitions

[Table 26-1](#) describes which releases of Veritas File System (VxFS) support specific SmartTier document type definitions (DTDs).

Table 26-1 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
6.0.1	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 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.

SmartTier 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 `fspadm` command or the GUI.

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

Note: Do not run the `fspadm` command simultaneously from different terminals.

The `lost+found` must exist before you can use the `fspadm` command.

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 `fspadm 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:

```
# fspadm 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 479.

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 `.__fsspadm_enforce.log` file if the file does not exist. The enforce operation appends the file if the file already exists. The `.__fsspadm_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 `fsspadm` 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 `fsspadm` 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 `fsspadm` command to sweep and relocate selectively. You can specify the `-T` option only if the policy uses the `Prefer` criteria for `IOTEMP`.

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

The following example uses the `fsspadm 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 the balance size 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.

For a multi-volume file system, you can specify the `compress` flag or the `uncompress` flag with the `<TO>` clause. The `compress` flag causes SmartTier to compress a file's extents while relocating the file to the tier specified by the `<DESTINATION>` element. SmartTier compresses the entire file and relocates the file to the destination tier, even if the file spans multiple tiers. The `uncompress` flag causes SmartTier to uncompress a file's extents while relocating the file to the tier specified by the `<DESTINATION>` element.

The following XML snippet specifies the `compress` flag:

```
<TO Flags="compress">
  <DESTINATION>
    <CLASS> tier4 </CLASS>
  </DESTINATION>
</TO>
```

The following XML snippet specifies the `uncompress` flag:

```
<TO Flags="uncompress">
  <DESTINATION>
    <CLASS> tier4 </CLASS>
  </DESTINATION>
</TO>
```

- `<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:

<code><ACCAGE></code>	This criterion is met when files are inactive for a designated period or during a designated period relative to the time at which the <code>fspadm enforce</code> command was issued.
<code><MODAGE></code>	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>fspadm enforce</code> command was issued.
<code><SIZE></code>	This criterion is met when files exceed or drop below a designated size or fall within a designated size range.
<code><IOTEMP></code>	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 <code><PERIOD></code> element prior to the time at which the <code>fspadm enforce</code> command was issued. See “Calculating I/O temperature and access temperature” on page 524.
<code><ACCESSTEMP></code>	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:

gt	The time of last access must be greater than the specified interval.
eq	The time of last access must be equal to the specified interval.
gteq	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.

lt	The time of last access must be less than the specified interval.
lteq	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
<code>GB</code>	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 524.

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 536.

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 701.

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 535.

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 536.

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:

$$\text{IOTEMP} = \frac{\text{(bytes of the file that are read in the PERIOD)}}{\text{(PERIOD in hours * size of the file in bytes)}}$$

The write and read/write IOTEMP are also computed accordingly.

The following formula computes the average read IOTEMP:

$$\text{Average IOTEMP} = \frac{\text{(bytes read of all active files in the last } h \text{ hours)}}{(h * \text{size of all the active files in bytes})}$$

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 ACCESTEMP

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 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>

```

```
</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 `fspadm 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 `fspadm 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.

The following example compresses while relocating all of the files from `tier2` with the extension `dbf` to `tier4` if the file was accessed over 30 days ago:

```

<SELECT Flags="Data">
    <PATTERN> *.dbf </PATTERN>
</SELECT>

<RELOCATE>
    <FROM>
        <SOURCE>
            <CLASS> tier2 </CLASS>
        </SOURCE>
    </FROM>
    <TO Flags="compress">
        <DESTINATION>
            <CLASS> tier4 </CLASS>
        </DESTINATION>
    </TO>
    <WHEN>
        <ACCAGE Units="days">
            <MIN Flags="gt">30</MIN>
        </ACCAGE>
    </WHEN>
</RELOCATE>

```

The following example uncompresses while relocating all of the files from `tier3` with the extension `dbf` to `tier1` if the file was accessed over 1 hour ago:

```

<SELECT Flags="Data">
  <PATTERN> *.dbf </PATTERN>
</SELECT>

<RELOCATE>
  <FROM>
    <SOURCE>
      <CLASS> tier3 </CLASS>
    </SOURCE>
  </FROM>
  <TO Flags="uncompress">
    <DESTINATION>
      <CLASS> tier1 </CLASS>
    </DESTINATION>
  </TO>
  <WHEN>
    <ACCAGE Units="hours">
      <MIN Flags="gt">1</MIN>
    </ACCAGE>
  </WHEN>
</RELOCATE>

```

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:

<FROM>	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 <WHEN> clause. No priority is associated with the ordering of placement classes 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 deleted. If a DELETE statement does not contain a <FROM> 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.

COMPRESS statement

The `COMPRESS` statement in a file placement policy rule specifies in-place file compression on multi-volume or single-volume file systems. The placement policy becomes assigned to the selected file, and allocation for the compressed extents is done from the same tier specified in the `<SOURCE>` element of the `<FROM>` clause. SmartTier performs in-place compression of the entire file, even if the file spans across multiple tiers.

Note: SmartTier does not schedule compression activity. If you did not integrate your Veritas Storage Foundation product with the Veritas Operations Manager (VOM), then you must automate compression activity by using techniques such as scheduling through cron jobs.

The following XML snippet illustrates the general form of the `COMPRESS` statement:

```
<COMPRESS>
  <FROM>
    <SOURCE>
      <CLASS> placement_class_name </CLASS>
    </SOURCE>
    <SOURCE> additional_placement_class_specifications
    </SOURCE>
  </FROM>
  <WHEN> compression_conditions </WHEN>
</COMPRESS>
```

A `COMPRESS` statement contains the following clauses:

<FROM>	An optional clause that contains a list of placement classes from whose volumes designated files should be compressed 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 compression. If a <code>COMPRESS</code> statement contains a <FROM> clause, VxFS only considers files that reside on volumes in placement classes specified in the clause for compression. If no <FROM> clause is present, qualifying files are compressed regardless of where the files reside.
<WHEN>	An optional clause that indicates the conditions under which files to which the rule applies should be compressed. 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 compressed. If a <code>COMPRESS</code> statement does not contain a <WHEN> clause, files to which the rule applies are compressed unconditionally. A <WHEN> clause may be included in a <code>COMPRESS</code> statement to specify that files should be compressed only if any or all of four types of criteria are met. Files can be specified for compression 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 <code>fspadm enforce</code> 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>fspadm 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>fspadm enforce</code> command was issued. See “Calculating I/O temperature and access temperature” on page 524.

<code><ACCESSTEMP></code>	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 `COMPRESS` 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 compression, 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:

gt	The time of last access must be greater than the specified interval.
eq	The time of last access must be equal to the specified interval.
gteq	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.

lt	The time of last access must be less than the specified interval.

`lteq` 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 compress files to which the rule applies that have been inactive for longer than the specified interval. Such a rule would typically be used to compress inactive files to less expensive storage tiers. Conversely, including `<MAX>` causes files accessed within the specified interval to be compressed. 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 compress files whose access age lies between the two.

The modification age compression 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 compression 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 compression criterion, `<SIZE>`, causes files to be compressed if the files are larger or smaller than the values specified in the `<MIN>` and `<MAX>` compression criteria, respectively, at the time that the `fsppadm enforce` command was issued. Specifying both criteria causes VxFS to schedule compression 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
<code>GB</code>	Gigabytes

Specifying the I/O temperature compression criterion

The I/O temperature compression criterion, `<IOTEMP>`, causes files to be compressed 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 524.

As with the other file compression 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 `fspadm 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 536.

For example, if you issued the `fspadm 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 `fspadm 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 701.

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 `fspadm 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 compression criterion. However, if the intention is to compress files as the I/O load on them decreases, then I/O temperature is more appropriate.

The case for upward compression is similar. If files that have been compressed to lower-tier storage devices due to infrequent access experience renewed application activity, then it may be appropriate to compress 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 compression criterion will cause files to be compressed that have been accessed even once during the interval. Conversely, a policy that uses I/O temperature with a `<MIN>` value will only compress 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 compressing files. The `Prefer` attribute can take two values: `low` or `high`. If you specify `low`, Veritas File System (VxFS) compresses the files with the lower I/O temperature before compressing the files with the higher I/O temperature. If you specify `high`, VxFS compresses the files with the higher I/O temperature before compressing 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 535.

Different `<PERIOD>` elements may be used in the `<IOTEMP>` and `<ACCESSTEMP>` criteria of different `COMPRESS` statements within the same policy.

The following placement policy snippet gives an example of the `Prefer` criteria:

```
<COMPRESS>
...
<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high">
    <MIN Flags="gteq"> 3.4 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</COMPRESS>
```

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 `COMPRESS` 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 `COMPRESS` 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 536.

The following placement policy snippet gives an example of the `Average` criteria:

```
<COMPRESS>
...
<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high" Average="*">
    <MIN Flags="gteq"> 1.5 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</COMPRESS>
```

In the snippet, VxFS compresss 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 compression 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">
```

COMPRESS statement examples

The following example compresses all of the files with the extension *dbf* on the multi-volume file system *tier2* that have not been accessed for last 30 days:

```
<SELECT Flags="Data">
  <PATTERN> *.dbf </PATTERN>
</SELECT>
```

```

<COMPRESS>
  <FROM>
    <SOURCE>
      <CLASS> tier2 </CLASS>
    </SOURCE>
  </FROM>
  <WHEN>
    <ACCAGE Units="days">
      <MIN Flags="gt">30</MIN>
    </ACCAGE>
  </WHEN>
</COMPRESS>

```

The files designated by the rule's `SELECT` statement that reside on volumes in placement class `tier2` at the time the `fsppadm enforce` command executes are compressed in place. Each instantiation of `fsppadm enforce` compresses files created since the last run on the `tier2` volumes.

The following example compresses all of the files with the extension `dbf` on a single volume if the file was not accessed for one minute.

```

<SELECT Flags="Data">
  <PATTERN> *.dbf </PATTERN>
</SELECT>

<COMPRESS>
  <WHEN>
    <ACCAGE Units="minutes">
      <MIN Flags="gt">1</MIN>
    </ACCAGE>
  </WHEN>
</COMPRESS>

```

No `<FROM>` clause is required for single volume. The files designated by the rule's `SELECT` statement at the time the `fsppadm enforce` command executes are compressed in place. Each instantiation of `fsppadm enforce` compresses files created since the last run on the volume.

The following example compresses all of the files on `tier3`:

```

<SELECT Flags="Data">
  <PATTERN> * </PATTERN>
</SELECT>

```

```

<COMPRESS>
  <FROM>
    <SOURCE>
      <CLASS> tier3 </CLASS>
    </SOURCE>
  </FROM>
</COMPRESS>

```

This rule compresses in place all files that reside on `tier3` at the time the `fspadm enforce` command executes.

UNCOMPRESS statement

The `UNCOMPRESS` statement in a file placement policy rule specifies in-place file uncompression on multi-volume and single-volume file systems. The placement policy becomes assigned to the selected file, and allocation for the uncompressed extents is done from the tier specified in the `<SOURCE>` element of the `<FROM>` clause.

If a file is partially compressed, then the file can be picked only for in-place compression, after which in next enforcement the file will be uncompressed before being relocated.

Note: SmartTier does not schedule uncompression activity. If you did not integrate your Veritas Storage Foundation product with the Veritas Operations Manager (VOM), then you must automate uncompression activity by using techniques such as scheduling through cron jobs.

The following XML snippet illustrates the general form of the `UNCOMPRESS` statement:

```

<UNCOMPRESS>
  <FROM>
    <SOURCE>
      <CLASS> placement_class_name </CLASS>
    </SOURCE>
    <SOURCE> additional_placement_class_specifications
    </SOURCE>
  </FROM>
  <WHEN> uncompression_conditions </WHEN>
</UNCOMPRESS>

```

A `UNCOMPRESS` statement contains the following clauses:

<FROM>	An optional clause that contains a list of placement classes from whose volumes designated files should be uncompressed 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 uncompression. If a <code>UNCOMPRESS</code> statement contains a <FROM> clause, VxFS only considers files that reside on volumes in placement classes specified in the clause for uncompression. If no <FROM> clause is present, qualifying files are uncompressed regardless of where the files reside.
<WHEN>	An optional clause that indicates the conditions under which files to which the rule applies should be uncompressed. 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 uncompressed. If a <code>UNCOMPRESS</code> statement does not contain a <WHEN> clause, files to which the rule applies are uncompressed unconditionally. A <WHEN> clause may be included in a <code>UNCOMPRESS</code> statement to specify that files should be uncompressed only if any or all of four types of criteria are met. Files can be specified for uncompression 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 <code>fspadm enforce</code> 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>fspadm 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>fspadm enforce</code> command was issued. See “Calculating I/O temperature and access temperature” on page 524.

<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 UNCOMPRESS 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 uncompression, 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:

gt	The time of last access must be greater than the specified interval.
eq	The time of last access must be equal to the specified interval.
gteq	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.

lt	The time of last access must be less than the specified interval.

lteq

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 uncompress files to which the rule applies that have been inactive for longer than the specified interval. Such a rule would typically be used to uncompress inactive files to less expensive storage tiers. Conversely, including `<MAX>` causes files accessed within the specified interval to be uncompressed. 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 uncompress files whose access age lies between the two.

The modification age uncompression 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 uncompression 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 uncompression criterion, `<SIZE>`, causes files to be uncompressed if the files are larger or smaller than the values specified in the `<MIN>` and `<MAX>` uncompression criteria, respectively, at the time that the `fsppadm enforce` command was issued. Specifying both criteria causes VxFS to schedule uncompression for files whose sizes lie between the two. Using the Units attribute, threshold file sizes may be specified in the following units:

bytes	Bytes
KB	Kilobytes
MB	Megabytes
GB	Gigabytes

Specifying the I/O temperature uncompression criterion

The I/O temperature uncompression criterion, `<IOTEMP>`, causes files to be uncompressed 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 524.

As with the other file uncompression 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 `fspadm 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 536.

For example, if you issued the `fspadm 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 `fspadm 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 701.

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 `fspadm 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 uncompression criterion. However, if the intention is to uncompress files as the I/O load on them decreases, then I/O temperature is more appropriate.

The case for upward uncompression is similar. If files that have been uncompressed to lower-tier storage devices due to infrequent access experience renewed application activity, then it may be appropriate to uncompress 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 uncompression criterion will cause files to be uncompressed that have been accessed even once during the interval. Conversely, a policy that uses I/O temperature with a `<MIN>` value will only uncompress 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 uncompressing files. The `Prefer` attribute can take two values: `low` or `high`. If you specify `low`, Veritas File System (VxFS) uncompresses the files with the lower I/O temperature before uncompressing the files with the higher I/O temperature. If you specify `high`, VxFS uncompresses the files with the higher I/O temperature before uncompressing 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 535.

Different `<PERIOD>` elements may be used in the `<IOTEMP>` and `<ACCESSTEMP>` criteria of different `UNCOMPRESS` statements within the same policy.

The following placement policy snippet gives an example of the `Prefer` criteria:

```
<UNCOMPRESS>
...
<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high">
    <MIN Flags="gteq"> 3.4 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</UNCOMPRESS>
```

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 `UNCOMPRESS` 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 `UNCOMPRESS` criteria specifies the 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 536.

The following placement policy snippet gives an example of the *Average* criteria:

```
<UNCOMPRESS>
...
<WHEN>
  <IOTEMP Type="nrbytes" Prefer="high" Average="*">
    <MIN Flags="gteq"> 1.5 </MIN>
    <PERIOD Units="hours"> 6 </PERIOD>
  </IOTEMP>
</WHEN>
</UNCOMPRESS>
```

In the snippet, VxFS uncompresss 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:

$$\text{IOTEMP} = \frac{\text{bytes of the file that are read in the PERIOD}}{(\text{PERIOD in hours} * \text{size of the file in bytes})}$$

The write and read/write IOTEMP are also computed accordingly.

The following formula computes the average read IOTEMP:

$$\text{Average IOTEMP} = \frac{\text{bytes read of all active files in the last } h \text{ hours}}{(h * \text{size of all the active files in bytes})}$$

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 uncompression 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">
```

UNCOMPRESS statement examples

The following example uncompresses in place all of the files with the extension `dbf` on the multi-volume file system `tier3` that have been accessed over 60 minutes ago:

```

<SELECT Flags="Data">
  <PATTERN> *.dbf </PATTERN>
</SELECT>

<UNCOMPRESS>
  <FROM>
    <SOURCE>
      <CLASS> tier3 </CLASS>
    </SOURCE>
  </FROM>
  <WHEN>
    <ACCAGE Units="minutes">
      <MIN Flags="gt">60</MIN>
    </ACCAGE>
  </WHEN>
</UNCOMPRESS>

```

The following example uncompresses in place all of the files with the extension `dbf` on a single volume that have been accessed over 1 minute ago:

```

<SELECT Flags="Data">
  <PATTERN> *.dbf </PATTERN>
</SELECT>

<UNCOMPRESS>
  <WHEN>
    <ACCAGE Units="minutes">
      <MIN Flags="gt">1</MIN>
    </ACCAGE>
  </WHEN>
</UNCOMPRESS>

```

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>` parameter.
- 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 701.

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 `fsppadm 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 535.
- Support of the `Prefer` attribute for the `<IOTEMP>` and `<ACCESSTEMP>` criteria
See [“Prefer mechanism with solid state disks”](#) on page 535.
- Provision of a mechanism to relocate based on average I/O activity
See [“Average I/O activity with solid state disks”](#) on page 536.
- 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 536.
- Quick identification of cold files

See [“Quick identification of cold files with solid state disks”](#) on page 537.

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 494.

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 495.

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 496.

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 `fsppadm` command with the `-C` option.
See [“Quick identification of cold files with solid state disks”](#) on page 537.
- 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 480.

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.

Sub-file relocation

The sub-file relocation functionality relocates the data ranges of the specified files to the specified target tier. Only one instance is allowed at a time on a given node for a given mount.

You can move sub-file data by using the `fsppadm subfilemove` command. The application using this framework calls the `fsppadm subfilemove` command periodically via some external scheduling mechanism at desired intervals, to effect relocations. The application might need to call `subfilemove` on each node of a cluster, in case of a cluster file system, if you want to distribute the load. The application also must arrange for initiating this relocation for new mounts and reboots, if the application needs sub-file relocations on those nodes or mounts.

In a cluster situation, since enforcement can happen from multiple nodes even if each node is scheduled to collect statistics at the same intervals, each node's persistence into the database can be slightly out of sync with each other on each node. Since enforcement should follow statistics collection, Symantec recommends that you schedule enforcements on each node with a few minutes of lag so that all nodes can complete the statistics synchronizing by that time. A lag time of 5 minutes suffices in most cases.

Note: You cannot use SmartTier to compress files while using the sub-file relocation functionality.

Moving sub-file data of files to specific target tiers

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

The following example moves a total of 32 MB in the file `test.dbf` from the offset 64 MB through 96 MB from its existing tier to `tier2`:

```
# cat /var/tmp/list
test.dbf 67108864 100663296 tier2
# fsppadm subfilemove -f /var/tmp/list /mount1
```

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 544.

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 667.

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 547.
See [“Complete disk failure mail messages”](#) on page 548.
See [“Modifying the behavior of hot-relocation”](#) on page 558.
- `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 27-1](#) shows the hot-relocation process in the case of the failure of a single subdisk of a RAID-5 volume.

Figure 27-1 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 558.

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 667.

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 547.

If the problem is not a cabling error, replace the disk.

See [“Removing and replacing disks”](#) on page 667.

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 551.

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 551.

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 552.

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 553.

See [“Making a disk available for hot-relocation use”](#) on page 554.

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 554.

After a successful relocation, remove and replace the failed disk.

See [“Removing and replacing disks”](#) on page 667.

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 667.

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 unrelocate 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 unrelocate 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`.

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.

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 327.

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 mount_point</i> all</code>
Set the deduplication schedule for a file system.	<code>fsdedupadm setschedule <i>time mount_point</i></code>

Functionality**Command syntax**

Initiate a deduplication dry run on a file system. `fsdedupadm dryrun [-o threshold=#] mount_point`

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 bytes 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 701.

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 564.

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.

You must start the `fsdedupschd` daemon before scheduling the task:

```
# /sbin/init.d/fsdedupschd start
```

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. You can perform a dry run only on a file system that has not been deduplicated previously.

See [“Enabling and disabling deduplication on a file system”](#) on page 564.

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 percent:

```
# 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 564.

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 564.

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)	Space_Shared (KB)
/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)	Space_Shared (KB)
/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)	Space_Shared(KB)
/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 573.

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, the backup program 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 29-1 describes how compressed files interact with other Veritas Storage Foundation commands and base operating system commands.

Table 29-1

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> command 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 reports the space savings due to compressed files. See the <code>fsadm_vxfs(1)</code> manual page.
<code>fsmap -p</code>	The <code>fsmap</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>fsmap(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 29-1 (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 29-2 describes how compressed files interact with other Veritas Storage Foundation features.

Table 29-2

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 29-2 (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 (SF), Veritas Storage Foundation and High Availability (SFHA), 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 the `winter_files` file.

```
SQL> select file_name from dba_data_files where  
tablespace_name = 'SUMMER_ITEM';
```

Store the result in the `summer_files` file.

- 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 the my_compress_files file.

- 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
```

Administering storage

- [Chapter 30. Managing volumes and disk groups](#)
- [Chapter 31. Rootability](#)
- [Chapter 32. Quotas](#)
- [Chapter 33. File Change Log](#)

Managing volumes and disk groups

This chapter includes the following topics:

- [Rules for determining the default disk group](#)
- [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 plexes and subdisks](#)
- [Decommissioning storage](#)

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 585.
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 585.
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.

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

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.

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

Use the following command to specify the name of the disk group that is aliased by `defaultdg`:

```
# vxctl defaultdg diskgroup
```

Where *diskgroup* is one of the following:

- A specified disk group name.
The specified disk group is not required to exist on the system.
- `bootdg`
Sets the default disk group to be the same as the currently defined system-wide boot disk group.
- `nodg`
Specifies that the default disk group is undefined.

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

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:

```
# vxdbg -g salesdg rmdisk salesdg04
# vxdbg -g mktdg adddisk mktdg02=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 595.

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 `vxldg` command to reorganize your disk groups.

The `vxldg` 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 30-1](#) shows the `move` operation.


Figure 30-1 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 30-2 shows the `split` operation.

Figure 30-2 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 30-3](#) shows the `join` operation.

Figure 30-3 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 `vxdg` 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 `vxdg` 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 622.
- 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.
See [“Setting the automatic recovery of volumes”](#) on page 629.
- 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:

```
# vxdbg [-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`:

```
# vxdbg 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`:

```
# vxdbg listmove mydg newdg mydg01
VxVM vxdbg ERROR V-5-2-4597 vxdbg listmove mydg newdg failed
VxVM vxdbg 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`:

```
# vxdbg -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 `vxddco` 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 30-4](#) shows some instances in which it is not 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 85.

Figure 30-4 Examples of disk groups that can and cannot be split



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:

```
# vxdbg [-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 `vxdbg` 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. `vxdbg` then checks again to ensure that the configuration has not changed since it performed the compatibility check. If the configuration has changed, `vxdbg` 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`:

```
# vxvg -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 vol1
```

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 595.

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  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           -           -           -
```

The following command removes disks `rootdg07` and `rootdg08` from `rootdg` to form a new disk group, `mydg`:

```
# vxvg -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	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	-	-	-	-
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	TUTILO	PUTILO
dg	mydg	mydg	-	-	-	-	-	-
dm	rootdg07	c1t99d0s2	-	17678493	-	-	-	-
dm	rootdg08	c1t100d0s2-	-	17678493	-	-	-	-

Joining disk groups

To remove all VxVM objects from an imported source disk group to an imported target disk group, use the following command:

```
# vxvg [-o override|verify] join sourcedg targetdg
```

See [“Moving objects between disk groups”](#) on page 595.

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`:

```
# vxdg 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	-	-	-

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.
<code>set</code>	Changes a task's modifiable parameters. Currently, there is only one modifiable parameter, <code>slow[=iodelay]</code>, which can be used to reduce the impact that copy operations have on system performance. If you specify <code>slow</code>, this introduces a delay between such operations with a default value for <code>iodelay</code> of 250 milliseconds. The larger <code>iodelay</code> value you specify, the slower the task progresses and the fewer system resources that it consumes in a given time. (The <code>vxplex</code>, <code>vxvol</code> and <code>vxrecover</code> commands also accept the <code>slow</code> 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 610.

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)
<code>span</code>	Concatenated
<code>stripe</code>	Striped

See [“Permitted layout transformations”](#) on page 606.

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 609.

Permitted relayout transformations

[Table 30-1](#) shows the supported relayout transformations for concatenated volumes.

Table 30-1 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 30-2](#) shows the supported relayout transformations for concatenated-mirror volumes.

Table 30-2 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.
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 30-3](#) shows the supported relayout transformations for RAID-5 volumes.

Table 30-3 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 30-4](#) shows the supported relayout transformations for mirror-concatenated volumes.

Table 30-4 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 30-5](#) shows the supported relayout transformations for mirrored-stripe volumes.

Table 30-5 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 30-6](#) shows the supported relayout transformations for unmirrored stripe and layered striped-mirror volumes.

Table 30-6 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 layout 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 layout 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 610.

See [“Controlling the progress of a relayout”](#) on page 610.

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 relayout

You can use the `vxtask` command to stop (pause) the relayout temporarily, or to cancel it (abort). If you specify a task tag to `vxassist` when you start the relayout, you can use this tag to specify the task to `vxtask`. For example, to pause the relayout operation that is tagged as `myconv`, enter:

```
# vxtask pause myconv
```

To resume the operation, use the `vxtask` command as follows:

```
# vxtask resume myconv
```

For relayout 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 relayout by specifying the `start` keyword to `vxrelayout`, as follows:

```
# vxrelayout -g mydg -o bg start vol04
```

If you use the `vxrelayout start` command to restart a relayout 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 relayout.

The `-o bg` option restarts the relayout in the background. You can also specify the `slow` and `iosize` option modifiers to control the speed of the relayout 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:

```
# vxrelayout -g mydg -o bg,slow=1000,iosize=10m start vol04
```

The default delay and region size values are 250 milliseconds and 1 MB respectively.

To reverse the direction of relayout operation that is stopped, specify the `reverse` keyword to `vxrelayout` as follows:

```
# vxrelayout -g mydg -o bg reverse vol04
```

This undoes changes made to the volume so far, and returns it to its original layout.

If you cancel a relayout 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 602.

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
none	Do not use SmartMove at all.
thinonly	Use SmartMove for thin aware LUNs only.
all	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 148.

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 `vol101-02` from the disk group `mydg`, use the following command:

```
# vxplex -g mydg -o rm dis vol101-02
```

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. A tag name must not be the same as the name of a disk in the disk group.

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`.

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`.

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 622.

[Table 30-7](#) 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 30-7 Disk group version assignments

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
6.0.1	180	CVM availability enhancements	20, 30, 40, 50, 60, 70, 80, 90, 110, 120, 130, 140, 150, 160, 170
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

Table 30-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
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
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

Table 30-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
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
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

Table 30-7 Disk group version assignments (*continued*)

VxVM release	Introduces disk group version	New features supported	Supports disk group versions
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
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 622.

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 6.0.1, 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:

```
# vxvg 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:

```
# vxvg 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.

Displaying disk group information

To display information on existing disk groups, enter the following command:

```
# vxdg list
NAME      STATE      ID
rootdg    enabled    730344554.1025.tweety
newdg     enabled    731118794.1213.tweety
```

To display more detailed information on a specific disk group, use the following command:

```
# vxdg list diskgroup
```

When you apply this command to a disk group named `mydg`, the output is similar to the following:

```
# vxdg list mydg

Group: mydg
dgid: 962910960.1025.bass
import-id: 0.1
flags:
version: 160
local-activation: read-write
alignment: 512 (bytes)
ssb: on
```

```
detach-policy: local
copies: nconfig=default nlog=default
config: seqno=0.1183 permlen=3448 free=3428 templen=12 loglen=522
config disk c0t10d0 copy 1 len=3448 state=clean online
config disk c0t11d0 copy 1 len=3448 state=clean online
log disk c0t10d0 copy 1 len=522
log disk c0t11d0 copy 1 len=522
```

To verify the disk group ID and name that is associated with a specific disk (for example, to import the disk group), use the following command:

```
# vxdisk -s list devicename
```

This command provides output that includes the following information for the specified disk. For example, output for disk `c0t12d0` as follows:

```
Disk:    c0t12d0
type:    simple
flags:   online ready private autoconfig autoimport imported
diskid:  963504891.1070.bass
dgroup:  newdg
dgid:    963504895.1075.bass
hostid:  bass
info:    privoffset=128
```

Displaying free space in a disk group

Before you add volumes and file systems to your system, make sure that you have enough free disk space to meet your needs.

To display free space in the system, use the following command:

```
# vxdg free
```

The following is example output:

GROUP	DISK	DEVICE	TAG	OFFSET	LENGTH	FLAGS
mydg	mydg01	c0t10d0	c0t10d0	0	4444228	-
mydg	mydg02	c0t11d0	c0t11d0	0	4443310	-
newdg	newdg01	c0t12d0	c0t12d0	0	4443310	-
newdg	newdg02	c0t13d0	c0t13d0	0	4443310	-
oradg	oradg01	c0t14d0	c0t14d0	0	4443310	-

To display free space for a disk group, use the following command:

```
# vxdg -g diskgroup free
```


where `-g diskgroup` optionally specifies a disk group.

For example, to display the free space in the disk group, `mydg`, use the following command:

```
# vxdg -g mydg free
```

The following example output shows the amount of free space in sectors:

DISK	DEVICE	TAG	OFFSET	LENGTH	FLAGS
mydg01	c0t10d0	c0t10d0	0	4444228	-
mydg02	c0t11d0	c0t11d0	0	4443310	-

Creating a disk group

You must associate a disk group with at least one disk. You can create a new disk group when you select **Add or initialize one or more disks** from the main menu of the `vxdiskadm` command to add disks to VxVM control. The disks to be added to a disk group must not belong to an existing disk group. A disk group name cannot include a period (.) character.

You can create a shared disk group.

You can also use the `vxdiskadd` command to create a new disk group. The command dialog is similar to that described for the `vxdiskadm` command.

In the following example, `c1t0d0` is the device name of a disk that is not currently assigned to a disk group.

```
# vxdiskadd c1t0d0
```

See [“Adding a disk to VxVM”](#) on page 273.

You can also create disk groups using the following `vxdg init` command:

```
# vxdg init diskgroup [cds=on|off] diskname=devicename
```

For example, to create a disk group named `mktdg` on device `c1t0d0s2`, enter the following:

```
# vxdg init mktdg mktdg01=c1t0d0s2
```

The disk that is specified by the device name, `c1t0d0s2`, must have been previously initialized with `vxdiskadd` or `vxdiskadm`. The disk must not currently belong to a disk group.

You can use the `cds` attribute with the `vxdg init` command to specify whether a new disk group is compatible with the Cross-platform Data Sharing (CDS) feature. Newly created disk groups are compatible with CDS by default (equivalent to

specifying `cds=on`). If you want to change this behavior, edit the file `/etc/default/vxdg` and set the attribute-value pair `cds=off` in this file before creating a new disk group.

You can also use the following command to set this attribute for a disk group:

```
# vxdg -g diskgroup set cds=on|off
```

Removing a disk from a disk group

Before you can remove the last disk from a disk group, you must disable the disk group.

See [“Disabling a disk group”](#) on page 655.

As an alternative to disabling the disk group, you can destroy it.

See [“Destroying a disk group”](#) on page 655.

If a disk contains no subdisks, you can remove it from its disk group with the following command:

```
# vxdg [-g diskgroup] rmdisk diskname
```

For example, to remove `mydg02` from the disk group `mydg`, enter the following:

```
# vxdg -g mydg rmdisk mydg02
```

If the disk has subdisks on it when you try to remove it, the following error message is displayed:

```
VxVM vxdg ERROR V-5-1-552 Disk diskname is used by one or more
subdisks
Use -k to remove device assignment.
```

Using the `-k` option lets you remove the disk even if it has subdisks.

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

Warning: Use of the `-k` option to `vxdg` can result in data loss.

After you remove the disk from its disk group, you can (optionally) remove it from VxVM control completely. Enter the following:

```
# vxdiskunsetup devicename
```

For example, to remove the disk `c1t0d0s2` from VxVM control, enter the following:

```
# vxdiskunsetup c1t0d0s2
```

You can remove a disk on which some subdisks of volumes are defined. For example, you can consolidate all the volumes onto one disk. If you use `vxdiskadm` to remove a disk, you can choose to move volumes off that disk. To do this, run `vxdiskadm` and select **Remove a disk** from the main menu.

If the disk is used by some volumes, this 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)
```

If you choose `y`, all volumes are moved off the disk, if possible. Some volumes may not be movable. The most common reasons why a volume may not be movable are as follows:

- There is not enough space on the remaining disks.
- Plexes or striped subdisks cannot be allocated on different disks from existing plexes or striped subdisks in the volume.

If `vxdiskadm` cannot move some volumes, you may need to remove some plexes from some disks to free more space before proceeding with the disk removal operation.

Deporting a disk group

Deporting a disk group disables access to a disk group that is enabled (imported) by the system. Deport a disk group if you intend to move the disks in a disk group to another system.

To deport a disk group

- 1 Stop all activity by applications to volumes that are configured in the disk group that is to be deported. Unmount file systems and shut down databases that are configured on the volumes.

If the disk group contains volumes that are in use (for example, by mounted file systems or databases), deportation fails.

- 2 To stop the volumes in the disk group, use the following command

```
# vxvol -g diskgroup stopall
```

- 3 From the `vxdiskadm` main menu, select **Remove access to (deport)** a disk group.

- 4 At prompt, enter the name of the disk group to be deported. In the following example it is `newdg`:

```
Enter name of disk group [<group>,list,q,?] (default: list)
newdg
```

- 5 At the following prompt, enter `y` if you intend to remove the disks in this disk group:

```
Disable (offline) the indicated disks? [y,n,q,?] (default: n) y
```

- 6 At the following prompt, press **Return** to continue with the operation:

```
Continue with operation? [y,n,q,?] (default: y)
```

After the disk group is deported, the `vxdiskadm` utility displays the following message:

```
VxVM INFO V-5-2-269 Removal of disk group newdg was
successful.
```

- 7 At the following prompt, indicate whether you want to disable another disk group (`y`) or return to the `vxdiskadm` main menu (`n`):

```
Disable another disk group? [y,n,q,?] (default: n)
```

You can use the following `vx dg` command to deport a disk group:

```
# vx dg deport diskgroup
```

Importing a disk group

Importing a disk group enables access by the system to a disk group. To move a disk group from one system to another, first disable (deport) the disk group on the original system, and then move the disk between systems and enable (import) the disk group.

By default, VxVM recovers and starts any disabled volumes in the disk group when you import the disk group. To prevent VxVM from recovering the disabled volumes, turn off the automatic recovery feature. For example, after importing the disk group, you may want to do some maintenance before starting the volumes.

See [“Setting the automatic recovery of volumes”](#) on page 629.

To import a disk group

- 1 To ensure that the disks in the deported disk group are online, use the following command:

```
# vxdisk -s list
```

- 2 From the `vxdiskadm` main menu, select Enable access to (import) a disk group.
- 3 At the following prompt, enter the name of the disk group to import (in this example, `newdg`):

```
Select disk group to import [<group>,list,q,?] (default: list)
newdg
```

When the import finishes, the `vxdiskadm` utility displays the following success message:

```
VxVM INFO V-5-2-374 The import of newdg was successful.
```

- 4 At the following prompt, indicate whether you want to import another disk group (y) or return to the `vxdiskadm` main menu (n):

```
Select another disk group? [y,n,q,?] (default: n)
```

You can also use the following `vx dg` command to import a disk group:

```
# vx dg import diskgroup
```

You can also import the disk group as a shared disk group.

Setting the automatic recovery of volumes

By default, VxVM recovers and starts any disabled volumes in the disk group when you import the disk group. To prevent VxVM from recovering the disabled volumes, turn off the automatic volume recovery. For example, after importing the disk group, you may want to do some maintenance before starting the volumes.

To turn off the automatic volume recovery feature

- ◆ Use the following `vx default` command to turn off automatic volume recovery.

```
# vx default set autostartvolumes off
```

Handling of minor number conflicts

If a minor conflict exists when a disk group is imported, VxVM automatically assigns a new base minor to the disk group, and reminors the volumes in the disk group, based on the new base minor. You do not need to run the `vxdbg reminor` command to resolve the minor conflicts.

To avoid any conflicts between shared and private disk groups, the minor numbers are divided into shared and private pools. VxVM allocates minor numbers of shared disk groups only from the shared pool, and VxVM allocates minor numbers of private disk groups only from the private pool. If you import a private disk group as a shared disk group or vice versa, the device minor numbers are re-allocated from the correct pool. The disk group is dynamically reminored.

By default, private minor numbers range from 0-32999, and shared minor numbers start from 33000. You can change the division if required. For example, you can set the range for shared minor numbers to start from a lower number. This range provides more minor numbers for shared disk groups and fewer minor numbers for private disk groups.

Normally the minor numbers in private and shared pools are sufficient, so there is no need to make changes to the division.

Note: To make the new division take effect, you must run `vxctl enable` or restart `vxconfigd` after the tunable is changed in the defaults file. The division on all the cluster nodes must be exactly the same, to prevent node failures for node join, volume creation, or disk group import operations.

To change the division between shared and private minor numbers

- 1 Add the tunable `sharedminorstart` to the defaults file `/etc/default/vxsf`. For example, to change the shared minor numbers so that the range starts from 20000, set the following line in the `/etc/default/vxsf` file.

```
sharedminorstart=20000
```

You cannot set the shared minor numbers to start at less than 1000. If `sharedminorstart` is set to values between 0 to 999, the division of private minor numbers and shared disk group minor numbers is set to 1000. The value of 0 disables dynamic renumbering.

- 2 Run the following command:

```
# vxctl enable
```

In certain scenarios, you may need to disable the division of between shared minor numbers and private minor numbers. For example, to prevent the device minor numbers from being changed when you upgrade from a previous release. In this case, disable the dynamic reminoring before you install the new VxVM package.

To disable the division between shared and private minor numbers

- 1 Set the tunable `sharedminorstart` in the defaults file `/etc/default/vxsf` to 0 (zero). Set the following line in the `/etc/default/vxsf` file.

```
sharedminorstart=0
```

- 2 Run the following command:

```
# vxctl enable
```

Moving disk groups between systems

An important feature of disk groups is that they can be moved between systems. If all disks in a disk group are moved from one system to another, then the disk group can be used by the second system. You do not have to re-specify the configuration.

To move a disk group between systems

- 1 Confirm that all disks in the diskgroup are visible on the target system. This may require masking and zoning changes.
- 2 On the source system, stop all volumes in the disk group, then deport (disable local access to) the disk group with the following command:

```
# vxctl deport diskgroup
```

- 3 Move all the disks to the target system and perform the steps necessary (system-dependent) for the target system and VxVM to recognize the new disks.

This can require a reboot, in which case the `vxconfigd` daemon is restarted and recognizes the new disks. If you do not reboot, use the command `vxctl enable` to restart the `vxconfigd` program so VxVM also recognizes the disks.

- 4 Import (enable local access to) the disk group on the target system with this command:

```
# vxvg import diskgroup
```

Warning: All disks in the disk group must be moved to the other system. If they are not moved, the import fails.

- 5 By default, VxVM enables and starts any disabled volumes after the disk group is imported.

See [“Setting the automatic recovery of volumes”](#) on page 629.

If the automatic volume recovery feature is turned off, start all volumes with the following command:

```
# vxrecover -g diskgroup -sb
```

You can also move disks from a system that has crashed. In this case, you cannot deport the disk group from the source system. When a disk group is created or imported on a system, that system writes a lock on all disks in the disk group.

Warning: The purpose of the lock is to ensure that SAN-accessed disks are not used by both systems at the same time. If two systems try to access the same disks at the same time, this must be managed using software such as the clustering functionality of VxVM. Otherwise, data and configuration information stored on the disk may be corrupted, and may become unusable.

Handling errors when importing disks

When you move disks from a system that has crashed or that failed to detect the group before the disk was moved, the locks stored on the disks remain and must be cleared. The system returns the following error message:

```
VxVM vxvg ERROR V-5-1-587 disk group groupname: import failed:  
Disk is in use by another host
```

The next message indicates that the disk group does not contain any valid disks (not that it does not contain any disks):

```
VxVM vxvg ERROR V-5-1-587 Disk group groupname: import failed:  
No valid disk found containing disk group
```


The disks may be considered invalid due to a mismatch between the host ID in their configuration copies and that stored in the `/etc/vx/volboot` file.

To clear locks on a specific set of devices, use the following command:

```
# vxdisk clearimport devicename ...
```

To clear the locks during import, use the following command:

```
# vxdg -C import diskgroup
```

Warning: Be careful when using the `vxdisk clearimport` or `vxdg -C import` command on systems that see the same disks via a SAN. Clearing the locks allows those disks to be accessed at the same time from multiple hosts and can result in corrupted data.

A disk group can be imported successfully if all the disks are accessible that were visible when the disk group was last imported successfully. However, sometimes you may need to specify the `-f` option to forcibly import a disk group if some disks are not available. If the `import` operation fails, an error message is displayed.

The following error message indicates a fatal error that requires hardware repair or the creation of a new disk group, and recovery of the disk group configuration and data:

```
VxVM vxdg ERROR V-5-1-587 Disk group groupname: import failed:
Disk group has no valid configuration copies
```

The following error message indicates a recoverable error.

```
VxVM vxdg ERROR V-5-1-587 Disk group groupname: import failed:
Disk for disk group not found
```

If some of the disks in the disk group have failed, you can force the disk group to be imported by specifying the `-f` option to the `vxdg import` command:

```
# vxdg -f import diskgroup
```

Warning: Be careful when using the `-f` option. It can cause the same disk group to be imported twice from different sets of disks. This can cause the disk group configuration to become inconsistent.

See [“Handling conflicting configuration copies”](#) on page 648.

As using the `-f` option to force the import of an incomplete disk group counts as a successful import, an incomplete disk group may be imported subsequently without this option being specified. This may not be what you expect.

You can also import the disk group as a shared disk group.

These operations can also be performed using the `vxdiskadm` utility. To deport a disk group using `vxdiskadm`, select `Remove access to (deport)` a disk group from the main menu. To import a disk group, select `Enable access to (import)` a disk group. The `vxdiskadm` import operation checks for host import locks and prompts to see if you want to clear any that are found. It also starts volumes in the disk group.

Reserving minor numbers for disk groups

A device minor number uniquely identifies some characteristic of a device to the device driver that controls that device. It is often used to identify some characteristic mode of an individual device, or to identify separate devices that are all under the control of a single controller. VxVM assigns unique device minor numbers to each object (volume, plex, subdisk, disk, or disk group) that it controls.

When you move a disk group between systems, it is possible for the minor numbers that it used on its previous system to coincide with those of objects known to VxVM on the new system. To get around this potential problem, you can allocate separate ranges of minor numbers for each disk group. VxVM uses the specified range of minor numbers when it creates volume objects from the disks in the disk group. This guarantees that each volume has the same minor number across reboots or reconfigurations. Disk groups may then be moved between machines without causing device number collisions.

VxVM chooses minor device numbers for objects created from this disk group starting at the base minor number `base_minor`. Minor numbers can range from this value up to 131,071. Try to leave a reasonable number of unallocated minor numbers near the top of this range to allow for temporary device number remapping in the event that a device minor number collision may still occur.

VxVM reserves the range of minor numbers from 0 to 999 for use with volumes in the boot disk group. For example, the `rootvol` volume is always assigned minor number 0.

If you do not specify the base of the minor number range for a disk group, VxVM chooses one at random. The number chosen is at least 1000, is a multiple of 1000, and yields a usable range of 1000 device numbers. The chosen number also does not overlap within a range of 1000 of any currently imported disk groups, and it does not overlap any currently allocated volume device numbers.

Note: The default policy ensures that a small number of disk groups can be merged successfully between a set of machines. However, where disk groups are merged automatically using failover mechanisms, select ranges that avoid overlap.

To view the base minor number for an existing disk group, use the `vxprint` command as shown in the following examples for the disk group, `mydg`:

```
# vxprint -l mydg | grep minors
minors: >=45000

# vxprint -g mydg -m | egrep base_minor
base_minor=45000
```

To set a base volume device minor number for a disk group that is being created, use the following command:

```
# vxdg init diskgroup minor=base_minor disk_access_name ...
```

For example, the following command creates the disk group, `newdg`, that includes the specified disks, and has a base minor number of 30000:

```
# vxdg init newdg minor=30000 c1d0t0s2 c1t1d0s2
```

If a disk group already exists, you can use the `vxdg reminor` command to change its base minor number:

```
# vxdg -g diskgroup reminor new_base_minor
```

For example, the following command changes the base minor number to 30000 for the disk group, `mydg`:

```
# vxdg -g mydg reminor 30000
```

If a volume is open, its old device number remains in effect until the system is rebooted or until the disk group is deported and re-imported. If you close the open volume, you can run `vxdg reminor` again to allow the renumbering to take effect without rebooting or re-importing.

An example of where it is necessary to change the base minor number is for a cluster-shareable disk group. The volumes in a shared disk group must have the same minor number on all the nodes. If there is a conflict between the minor numbers when a node attempts to join the cluster, the join fails. You can use the `reminor` operation on the nodes that are in the cluster to resolve the conflict. In a cluster where more than one node is joined, use a base minor number which does not conflict on any node.

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

See [“Handling of minor number conflicts”](#) on page 630.

Compatibility of disk groups between platforms

For disk groups that support the Cross-platform Data Sharing (CDS) feature, the upper limit on the minor number range is restricted on AIX, HP-UX, Linux (with a 2.6 or later kernel) and Solaris to 65,535 to ensure portability between these operating systems.

On a Linux platform with a pre-2.6 kernel, the number of minor numbers per major number is limited to 256 with a base of 0. This has the effect of limiting the number of volumes and disks that can be supported system-wide to a smaller value than is allowed on other operating system platforms. The number of disks that are supported by a pre-2.6 Linux kernel is typically limited to a few hundred. With the extended major numbering scheme that was implemented in VxVM 4.0 on Linux, a maximum of 4079 volumes could be configured, provided that a contiguous block of 15 extended major numbers was available.

VxVM 4.1 and later releases run on a 2.6 version Linux kernel, which increases the number of minor devices that are configurable from 256 to 65,536 per major device number. This allows a large number of volumes and disk devices to be configured on a system. The theoretical limit on the number of DMP and volume devices that can be configured on such a system are 65,536 and 1,048,576 respectively. However, in practice, the number of VxVM devices that can be configured in a single disk group is limited by the size of the private region.

When a CDS-compatible disk group is imported on a Linux system with a pre-2.6 kernel, VxVM attempts to reassign the minor numbers of the volumes, and fails if this is not possible.

To help ensure that a CDS-compatible disk group is portable between operating systems, including Linux with a pre-2.6 kernel, use the following command to set the `maxdev` attribute on the disk group:

```
# vxdbg -g diskgroup set maxdev=4079
```

Note: Such a disk group may still not be importable by VxVM 4.0 on Linux with a pre-2.6 kernel if it would increase the number of minor numbers on the system that are assigned to volumes to more than 4079, or if the number of available extended major numbers is smaller than 15.

You can use the following command to discover the maximum number of volumes that are supported by VxVM on a Linux host:

```
# cat /proc/sys/vxvm/vxio/vol_max_volumes  
4079
```

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

Handling cloned disks with duplicated identifiers

A disk may be copied by creating a hardware snapshot (such as an EMC BCV™ or Hitachi ShadowCopy™) or clone, by using `dd` or a similar command to replicate the disk, or by building a new LUN from the space that was previously used by a deleted LUN. To avoid the duplicate disk ID condition, the default action of VxVM is to prevent such duplicated disks from being imported.

Advanced disk arrays provide hardware tools that you can use to create clones of existing disks outside the control of VxVM. For example, these disks may have been created as hardware snapshots or mirrors of existing disks in a disk group. As a result, the VxVM private region is also duplicated on the cloned disk. When the disk group containing the original disk is subsequently imported, VxVM detects multiple disks that have the same disk identifier that is defined in the private region. In releases prior to 5.0, if VxVM could not determine which disk was the original, it would not import such disks into the disk group. The duplicated disks would have to be re-initialized before they could be imported.

From release 5.0, a unique disk identifier (UDID) is added to the disk's private region when the disk is initialized or when the disk is imported into a disk group (if this identifier does not already exist). Whenever a disk is brought online, the current UDID value that is known to the Device Discovery Layer (DDL) is compared with the UDID that is set in the disk's private region. If the UDID values do not match, the `udid_mismatch` flag is set on the disk. This flag can be viewed with the `vxdisk list` command. This allows a LUN snapshot to be imported on the same host as the original LUN. It also allows multiple snapshots of the same LUN to be simultaneously imported on a single server, which can be useful for off-host backup and processing.

A new set of `vxdisk` and `vxvg` operations are provided to handle such disks; either by writing the DDL value of the UDID to a disk's private region, or by tagging a disk and specifying that it is a cloned disk to the `vxvg import` operation.

The following is sample output from the `vxdisk list` command showing that disks `c2t66d0s2`, `c2t67d0s2` and `c2t68d0s2` are marked with the `udid_mismatch` flag:

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	auto:cdsdisk	-	-	online
c0t1d0s2	auto:cdsdisk	-	-	online
.				
.				

```

c2t64d0s2    auto:cdsdisk - - online
c2t65d0s2    auto:cdsdisk - - online
c2t66d0s2    auto:cdsdisk - - online udid_mismatch
c2t67d0s2    auto:cdsdisk - - online udid_mismatch
c2t68d0s2    auto:cdsdisk - - online udid_mismatch

```

Writing a new UDID to a disk

You can use the following command to update the unique disk identifier (UDID) for one or more disks. This is useful when building a new LUN from space previously used by a deleted LUN, for example.

```
# vxdisk [-f] [-g diskgroup] updateudid disk ...
```

This command uses the current value of the UDID that is stored in the Device Discovery Layer (DDL) database to correct the value in the private region. The `-f` option must be specified if VxVM has not set the `udid_mismatch` flag for a disk.

For example, the following command updates the UDIDs for the disks `c2t66d0s2` and `c2t67d0s2`:

```
# vxdisk updateudid c2t66d0s2 c2t67d0s2
```

Importing a disk group containing cloned disks

By default, disks on which the `udid_mismatch` flag or the `clone_disk` flag has been set are not imported by the `vxvg import` command unless all disks in the disk group have at least one of these flags set, and no two of the disks have the same UDID. You can then import the cloned disks by specifying the `-o useclonedev=on` option to the `vxvg import` command, as shown in this example:

```
# vxvg -o useclonedev=on [-o updateid] import mydg
```

This form of the command allows only cloned disks to be imported. All non-cloned disks remain unimported.

If the `clone_disk` flag is set on a disk, this indicates the disk was previously imported into a disk group with the `udid_mismatch` flag set.

The `-o updateid` option can be specified to write new identification attributes to the disks, and to set the `clone_disk` flag on the disks. (The `vxdisk set clone=on` command can also be used to set the flag.) However, the import fails if multiple copies of one or more cloned disks exist. In this case, you can use the following command to tag all the disks in the disk group that are to be imported:

```
# vxdisk [-g diskgroup] settag tagname disk ...
```

where *tagname* is a string of up to 128 characters, not including spaces or tabs.

For example, the following command sets the tag, `my_tagged_disks`, on several disks that are to be imported together:

```
# vxdisk settag my_tagged_disks c2t66d0s2 c2t67d0s2
```

Alternatively, you can update the UDIDs of the cloned disks.

See [“Writing a new UDID to a disk”](#) on page 638.

To check which disks are tagged, use the `vxdisk listtag` command:

```
# vxdisk listtag
```

DEVICE	NAME	VALUE
c0t0d0s2	-	-
c0t1d0s2	-	-
.		
.		
.		
c2t64d0s2	my_tagged_disks	-
c2t65d0s2	my_tagged_disks	-
c2t66d0s2	my_tagged_disks	-
c2t67d0s2	my_tagged_disks	-
c2t68d0s2	-	-

The configuration database in a VM disk’s private region contains persistent configuration data (or metadata) about the objects in a disk group. This database is consulted by VxVM when the disk group is imported. At least one of the cloned disks that are being imported must contain a copy of the current configuration database in its private region.

You can use the following command to ensure that a copy of the metadata is placed on a disk, regardless of the placement policy for the disk group:

```
# vxdisk [-g diskgroup] set disk keepmeta=always
```

Alternatively, use the following command to place a copy of the configuration database and kernel log on all disks in a disk group that share a given tag:

```
# vxdg [-g diskgroup] set tagmeta=on tag=tagname nconfig=all \
nlog=all
```

To check which disks in a disk group contain copies of this configuration information, use the `vxdg listmeta` command:

```
# vxdg [-q] listmeta diskgroup
```

The `-q` option can be specified to suppress detailed configuration information from being displayed.

The tagged disks in the disk group may be imported by specifying the tag to the `vxldg import` command in addition to the `-o useclonedev=on` option:

```
# vxldg -o useclonedev=on -o tag=my_tagged_disks import mydg
```

If you have already imported the non-cloned disks in a disk group, you can use the `-n` and `-t` option to specify a temporary name for the disk group containing the cloned disks:

```
# vxldg -t -n clonedg -o useclonedev=on -o tag=my_tagged_disks \  
import mydg
```

See [“Renaming a disk group”](#) on page 646.

To remove a tag from a disk, use the `vxldisk rmtag` command, as shown in the following example:

```
# vxldisk rmtag tag=my_tagged_disks c2t67d0s2
```

See the `vxldisk(1M)` and `vxldg(1M)` manual pages.

Sample cases of operations on cloned disks

The following sections contain examples of operations that can be used with cloned disks:

See [“Enabling configuration database copies on tagged disks”](#) on page 640.

See [“Importing cloned disks without tags”](#) on page 641.

See [“Importing cloned disks with tags”](#) on page 643.

Enabling configuration database copies on tagged disks

In this example, the following commands have been used to tag some of the disks in an Hitachi TagmaStore array:

```
# vxldisk settag TagmaStore-USP0_28 t1=v1  
# vxldisk settag TagmaStore-USP0_28 t2=v2  
# vxldisk settag TagmaStore-USP0_24 t2=v2  
# vxldisk settag TagmaStore-USP0_25 t1=v1
```

These tags can be viewed by using the `vxldisk listtag` command:

```
# vxldisk listtag
```


DEVICE	NAME	VALUE
TagmaStore-USP0_24	t2	v2
TagmaStore-USP0_25	t1	v1
TagmaStore-USP0_28	t1	v1
TagmaStore-USP0_28	t2	v2

The following command ensures that configuration database copies and kernel log copies are maintained for all disks in the disk group `mydg` that are tagged as `t1`:

```
# vxdbg -g mydg set tagmeta=on tag=t1 nconfig=all nlog=all
```

The disks for which such metadata is maintained can be seen by using this command:

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_10	auto:cdsdisk	-	-	online
TagmaStore-USP0_24	auto:cdsdisk	mydg02	mydg	online
TagmaStore-USP0_25	auto:cdsdisk	mydg03	mydg	online tagmeta
TagmaStore-USP0_26	auto:cdsdisk	-	-	online
TagmaStore-USP0_27	auto:cdsdisk	-	-	online
TagmaStore-USP0_28	auto:cdsdisk	mydg01	mydg	online tagmeta

Alternatively, the following command can be used to ensure that a copy of the metadata is kept with a disk:

```
# vxdisk set TagmaStore-USP0_25 keepmeta=always
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_10	auto:cdsdisk	-	-	online
TagmaStore-USP0_22	auto:cdsdisk	-	-	online
TagmaStore-USP0_23	auto:cdsdisk	-	-	online
TagmaStore-USP0_24	auto:cdsdisk	mydg02	mydg	online
TagmaStore-USP0_25	auto:cdsdisk	mydg03	mydg	online keepmeta
TagmaStore-USP0_28	auto:cdsdisk	mydg01	mydg	online

Importing cloned disks without tags

In the first example, cloned disks (ShadowImage™ devices) from an Hitachi TagmaStore array will be imported. The primary (non-cloned) disks, `mydg01`, `mydg02` and `mydg03`, are already imported, and the cloned disks are not tagged.

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_3	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_23	auto:cdsdisk	mydg02	mydg	online
TagmaStore-USP0_25	auto:cdsdisk	mydg03	mydg	online
TagmaStore-USP0_30	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_31	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_32	auto:cdsdisk	mydg01	mydg	online

To import the cloned disks, they must be assigned a new disk group name, and their UDIDs must be updated:

```
# vxvxdg -n snapdg -o useclonedev=on -o updateid import mydg
# vxvxdisk -o allldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_3	auto:cdsdisk	mydg03	snapdg	online clone_disk
TagmaStore-USP0_23	auto:cdsdisk	mydg02	mydg	online
TagmaStore-USP0_25	auto:cdsdisk	mydg03	mydg	online
TagmaStore-USP0_30	auto:cdsdisk	mydg02	snapdg	online clone_disk
TagmaStore-USP0_31	auto:cdsdisk	mydg01	snapdg	online clone_disk
TagmaStore-USP0_32	auto:cdsdisk	mydg01	mydg	online

Note that the state of the imported cloned disks has changed from `online udid_mismatch` to `online clone_disk`.

In the next example, none of the disks (neither cloned nor non-cloned) have been imported:

```
# vxvxdisk -o allldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_3	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_23	auto:cdsdisk	-	(mydg)	online
TagmaStore-USP0_25	auto:cdsdisk	-	(mydg)	online
TagmaStore-USP0_30	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_31	auto:cdsdisk	-	(mydg)	online udid_mismatch
TagmaStore-USP0_32	auto:cdsdisk	-	(mydg)	online

To import only the cloned disks into the `mydg` disk group:

```
# vxvxdg -o useclonedev=on -o updateid import mydg
# vxvxdisk -o allldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
TagmaStore-USP0_3	auto:cdsdisk	mydg03	mydg	online clone_disk

```

TagmaStore-USP0_23 auto:cdsdisk - (mydg) online
TagmaStore-USP0_25 auto:cdsdisk - (mydg) online
TagmaStore-USP0_30 auto:cdsdisk mydg02 mydg online clone_disk
TagmaStore-USP0_31 auto:cdsdisk mydg01 mydg online clone_disk
TagmaStore-USP0_32 auto:cdsdisk - (mydg) online

```

In the next example, a cloned disk (BCV device) from an EMC Symmetrix DMX array is to be imported. Before the cloned disk, EMC0_27, has been split off from the disk group, the `vxdisk list` command shows that it is in the `error udid_mismatch` state:

```

# vxdisk -o alldgs list
DEVICE          TYPE          DISK      GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1  mydg   online
EMC0_27         auto          -        -      error udid_mismatch

```

The `symmir` command is used to split off the BCV device:

```
# /usr/symcli/bin/symmir -g mydg split DEV001
```

After updating VxVM's information about the disk by running the `vxdisk scandisks` command, the cloned disk is in the `online udid_mismatch` state:

```

# vxdisk -o alldgs list
DEVICE          TYPE          DISK      GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1  mydg   online
EMC0_27         auto:cdsdisk -        -      online udid_mismatch

```

The following command imports the cloned disk into the new disk group `newdg`, and updates the disk's UDID:

```
# vxdg -n newdg -o useclonedev=on -o updateid import mydg
```

The state of the cloned disk is now shown as `online clone_disk`:

```

# vxdisk -o alldgs list
DEVICE          TYPE          DISK      GROUP  STATUS
EMC0_1          auto:cdsdisk EMC0_1  mydg   online
EMC0_27         auto:cdsdisk EMC0_1  newdg  online clone_disk

```

Importing cloned disks with tags

In this example, cloned disks (BCV devices) from an EMC Symmetrix DMX array will be imported. The primary (non-cloned) disks, `mydg01`, `mydg02` and `mydg03`, are already imported, and the cloned disks with the tag `t1` are to be imported.

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
EMC0_4	auto:cdsdisk	mydg01	mydg	online
EMC0_6	auto:cdsdisk	mydg02	mydg	online
EMC0_8	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_15	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_18	auto:cdsdisk	mydg03	mydg	online
EMC0_24	auto:cdsdisk	-	(mydg)	online udid_mismatch

The disks are tagged as follows:

```
# vxdisk listtag
```

DEVICE	NAME	VALUE
EMC0_4	t2	v2
EMC0_4	t1	v1
EMC0_6	t2	v2
EMC0_8	t1	v1
EMC0_15	t2	v2
EMC0_18	t1	v1
EMC0_24	t1	v1
EMC0_24	t2	v2

To import the cloned disks that are tagged as `t1`, they must be assigned a new disk group name, and their UDIDs must be updated:

```
# vxdg -n bcvdg -o useclonedev=on -o tag=t1 -o updateid import mydg
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
EMC0_4	auto:cdsdisk	mydg01	mydg	online
EMC0_6	auto:cdsdisk	mydg02	mydg	online
EMC0_8	auto:cdsdisk	mydg03	bcvdg	online clone_disk
EMC0_15	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_18	auto:cdsdisk	mydg03	mydg	online
EMC0_24	auto:cdsdisk	mydg01	bcvdg	online clone_disk

As the cloned disk `EMC0_15` is not tagged as `t1`, it is not imported. Note that the state of the imported cloned disks has changed from `online udid_mismatch` to `online clone_disk`.

By default, the state of imported cloned disks is shown as `online clone_disk`. This can be removed by using the `vxdisk set` command as shown here:

```
# vxdisk set EMC0_8 clone=off
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
EMC0_4	auto:cdsdisk	mydg01	mydg	online
EMC0_6	auto:cdsdisk	mydg02	mydg	online
EMC0_8	auto:cdsdisk	mydg03	bcvdg	online
EMC0_15	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_18	auto:cdsdisk	mydg03	mydg	online
EMC0_24	auto:cdsdisk	mydg01	bcvdg	online clone_disk

In the next example, none of the disks (neither cloned nor non-cloned) have been imported:

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
EMC0_4	auto:cdsdisk	-	(mydg)	online
EMC0_6	auto:cdsdisk	-	(mydg)	online
EMC0_8	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_15	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_18	auto:cdsdisk	-	(mydg)	online
EMC0_24	auto:cdsdisk	-	(mydg)	online udid_mismatch

To import only the cloned disks that have been tagged as `t1` into the `mydg` disk group:

```
# vxdg -o useclonedev=on -o tag=t1 -o updateid import mydg
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
EMC0_4	auto:cdsdisk	-	(mydg)	online
EMC0_6	auto:cdsdisk	-	(mydg)	online
EMC0_8	auto:cdsdisk	mydg03	mydg	online clone_disk
EMC0_15	auto:cdsdisk	-	(mydg)	online udid_mismatch
EMC0_18	auto:cdsdisk	-	(mydg)	online
EMC0_24	auto:cdsdisk	mydg01	mydg	online clone_disk

As in the previous example, the cloned disk `EMC0_15` is not tagged as `t1`, and so it is not imported.

Considerations when using EMC CLARiiON SNAPSHOT LUNs

If you need to import the Snapshot LUN of a primary LUN to the same host as the original LUN, be aware of the following limitation.

If you are using Enclosure-based naming (EBN) with the Array Volume id (AVID) enabled, turn off name persistence during device discovery before importing the snapshot LUN to the original host.

To turn off name persistence, use the following command:

```
# vxddladm set namingscheme=ebn persistence=no use_avid=yes
```

After DDL recognizes the LUN, turn on name persistence using the following command:

```
# vxddladm set namingscheme=ebn persistence=yes use_avid=yes
```

Renaming a disk group

Only one disk group of a given name can exist per system. It is not possible to import or deport a disk group when the target system already has a disk group of the same name. To avoid this problem, VxVM allows you to rename a disk group during import or deport.

To rename a disk group during import, use the following command:

```
# vxdg [-t] -n newdg import diskgroup
```

If the `-t` option is included, the import is temporary and does not persist across reboots. In this case, the stored name of the disk group remains unchanged on its original host, but the disk group is known by the name specified by *newdg* to the importing host. If the `-t` option is not used, the name change is permanent.

For example, this command temporarily renames the disk group, *mydg*, as *mytempdg* on import:

```
# vxdg -t -n mytempdg import mydg
```

To rename a disk group during deport, use the following command:

```
# vxdg [-h hostname] -n newdg deport diskgroup
```

When renaming on deport, you can specify the `-h hostname` option to assign a lock to an alternate host. This ensures that the disk group is automatically imported when the alternate host reboots.

For example, this command renames the disk group, *mydg*, as *myexdg*, and departs it to the host, *jingo*:

```
# vxdg -h jingo -n myexdg deport mydg
```

You cannot use this method to rename the boot disk group because it contains volumes that are in use by mounted file systems (such as `/`). To rename the boot disk group, you must first unmirror and unencapsulate the root disk, and then re-encapsulate and remirror the root disk in a different disk group. This disk group becomes the new boot disk group.

To temporarily move the boot disk group, `bootdg`, from one host to another (for repair work on the root volume, for example) and then move it back

- 1 On the original host, identify the disk group ID of the `bootdg` disk group to be imported with the following command:

```
# vxdisk -g bootdg -s list
```

```
dgname: rootdg
```

```
dgid: 774226267.1025.tweety
```

In this example, the administrator has chosen to name the boot disk group as `rootdg`. The ID of this disk group is `774226267.1025.tweety`.

This procedure assumes that all the disks in the boot disk group are accessible by both hosts.

- 2 Shut down the original host.

- 3 On the importing host, import and rename the `rootdg` disk group with this command:

```
# vxdg -tC -n newdg import diskgroup
```

The `-t` option indicates a temporary import name, and the `-C` option clears import locks. The `-n` option specifies an alternate name for the `rootdg` being imported so that it does not conflict with the existing `rootdg`. `diskgroup` is the disk group ID of the disk group being imported (for example, `774226267.1025.tweety`).

If a reboot or crash occurs at this point, the temporarily imported disk group becomes unimported and requires a reimport.

- 4 After the necessary work has been done on the imported disk group, deport it back to its original host with this command:

```
# vxdg -h hostname deport diskgroup
```

Here `hostname` is the name of the system whose `rootdg` is being returned (the system name can be confirmed with the command `uname -n`).

This command removes the imported disk group from the importing host and returns locks to its original host. The original host can then automatically import its boot disk group at the next reboot.

Handling conflicting configuration copies

If an incomplete disk group is imported on several different systems, this can create inconsistencies in the disk group configuration copies that you may need to resolve manually. This section and following sections describe how such a condition can occur, and how to correct it. (When the condition occurs in a cluster that has been split, it is usually referred to as a serial split brain condition).


Example of a serial split brain condition in a cluster

This section presents an example of how a serial split brain condition might occur for a shared disk group in a cluster. Conflicts between configuration copies can also occur for private disk groups in clustered and non-clustered configurations where the disk groups have been partially imported on different systems.

A campus cluster (also known as a stretch cluster or remote mirror configuration) typically consists of a 2-node cluster where each component (server, switch and storage) of the cluster exists in a separate building.

Figure 30-5 shows a 2-node cluster with node 0, a fibre channel switch and disk enclosure `enc0` in building A, and node 1, another switch and enclosure `enc1` in building B.

Figure 30-5 Typical arrangement of a 2-node campus cluster



The fibre channel connectivity is multiply redundant to implement redundant-loop access between each node and each enclosure. As usual, the two nodes are also linked by a redundant private network.

A serial split brain condition typically arises in a cluster when a private (non-shared) disk group is imported on Node 0 with Node 1 configured as the failover node.

If the network connections between the nodes are severed, both nodes think that the other node has died. (This is the usual cause of the split brain condition in clusters). If a disk group is spread across both enclosure `enc0` and `enc1`, each portion loses connectivity to the other portion of the disk group. Node 0 continues to update to the disks in the portion of the disk group that it can access. Node 1, operating as the failover node, imports the other portion of the disk group (with the `-f` option set), and starts updating the disks that it can see.

When the network links are restored, attempting to reattach the missing disks to the disk group on Node 0, or to re-import the entire disk group on either node, fails.


VxVM increments the serial ID in the disk media record of each imported disk in all the disk group configuration databases on those disks, and also in the private region of each imported disk. The value that is stored in the configuration database represents the serial ID that the disk group expects a disk to have. The serial ID that is stored in a disk's private region is considered to be its actual value. VxVM detects the serial split brain when the actual serial ID of the disks that are being attached mismatches with the serial ID in the disk group configuration database of the imported disk group.

If some disks went missing from the disk group (due to physical disconnection or power failure) and those disks were imported by another host, the serial IDs for the disks in their copies of the configuration database, and also in each disk's private region, are updated separately on that host. When the disks are subsequently re-imported into the original shared disk group, the actual serial IDs on the disks do not agree with the expected values from the configuration copies on other disks in the disk group.

Depending on what happened to the different portions of the split disk group, there are two possibilities for resolving inconsistencies between the configuration databases:


- If the other disks in the disk group were not imported on another host, VxVM resolves the conflicting values of the serial IDs by using the version of the configuration database from the disk with the greatest value for the updated ID (shown as `update_id` in the output from the `vxvg list diskgroup` command). [Figure 30-6](#) shows an example of a serial split brain condition that can be resolved automatically by VxVM.

Figure 30-6 Example of a serial split brain condition that can be resolved automatically



- If the other disks were also imported on another host, no disk can be considered to have a definitive copy of the configuration database.
[Figure 30-7](#) shows an example of a true serial split brain condition that cannot be resolved automatically by VxVM.

Figure 30-7 Example of a true serial split brain condition that cannot be resolved automatically



In this case, the disk group import fails, and the `vxvg` utility outputs error messages similar to the following before exiting:

```
VxVM vxconfigd NOTICE V-5-0-33 Split Brain. da id is 0.1, while dm id
is 0.0 for DM mydg01
VxVM vxvg ERROR V-5-1-587 Disk group newdg: import failed: Serial
Split Brain detected. Run vxsplitlines
```

The import does not succeed even if you specify the `-f` flag to `vxvg`.

Although it is usually possible to resolve this conflict by choosing the version of the configuration database with the highest valued configuration ID (shown as the value of `seqno` in the output from the `vxvg list diskgroup| grep config` command), this may not be the correct thing to do in all circumstances.

See [“Correcting conflicting configuration information”](#) on page 653.

Correcting conflicting configuration information

To resolve conflicting configuration information, you must decide which disk contains the correct version of the disk group configuration database. To assist you in doing this, you can run the `vxsplitlines` command to show the actual serial ID on each disk in the disk group and the serial ID that was expected from the configuration database. For each disk, the command also shows the `vx dg` command that you must run to select the configuration database copy on that disk as being the definitive copy to use for importing the disk group.

Note: The disk group must have a version number of at least 110.

The following is sample output from running `vxsplitlines` on the disk group `newdg`:

```
# vxsplitlines -v -g newdg
```

```
VxVM. vxsplitlines NOTICE V-0-0-0 There are 2 pools
All the disks in the first pool have the same config copies
All the disks in the second pool may not have the same config copies
```

To see the configuration copy from a disk, enter the following command:

```
# /etc/vx/diag.d/vxprivutil dumpconfig private path
```

To import the disk group with the configuration copy from a disk, enter the following command:

```
# /usr/sbin/vxdg (-s) -o selectcp=diskid import newdg
```

```
Pool 0
DEVICE DISK DISK ID DISK PRIVATE PATH
newdg1 sdp 1215378871.300.vm28501x13 /dev/vx/rdmp/sdp5
newdg2 sdq 1215378871.300.vm28501x13 /dev/vx/rdmp/sdp5
```

```
Pool 1
DEVICE DISK DISK ID DISK PRIVATE PATH
newdg3 sdo 1215378871.294.vm28501x13 /dev/vx/rdmp/sdo5
```

If you do not specify the `-v` option, the command has the following output:

```
# vxsplitlines -g mydg listssbinfo
```

```
VxVM vxdg listssbinfo NOTICE V-0-0-0 There are 2 pools
All the disks in the first pool have the same config copies
All the disks in the second pool may not have the same config copies
```

Number of disks in the first pool: 1
Number of disks in the second pool: 1

To import the disk group with the configuration copy from the first pool, enter the following command:

```
# /usr/sbin/vxdg (-s) -o selectcp=1221451925.395.vm28501x13 import mydg
```

To import the disk group with the configuration copy from the second pool, enter the following command:

```
# /usr/sbin/vxdg (-s) -o selectcp=1221451927.401.vm28501x13 import mydg
```

In this example, the disk group has four disks, and is split so that two disks appear to be on each side of the split.

You can specify the `-c` option to `vxsplitlines` to print detailed information about each of the disk IDs from the configuration copy on a disk specified by its disk access name:

```
# vxsplitlines -g newdg -c c2t6d0s2
DANAME(DMNAME) || Actual SSB      || Expected SSB
c2t5d0s2( c2t5d0s2 ) || 0.1          || 0.0 ssb ids don't match
c2t6d0s2( c2t6d0s2 ) || 0.1          || 0.1 ssb ids match
c2t7d0s2( c2t7d0s2 ) || 0.1          || 0.1 ssb ids match
c2t8d0s2( c2t8d0s2 ) || 0.1          || 0.0 ssb ids don't match
```

Please note that even though some disks ssb ids might match that does not necessarily mean that those disks' config copies have all the changes. From some other configuration copies, those disks' ssb ids might not match. To see the configuration from this disk, run

```
/etc/vx/diag.d/vxprivutil dumpconfig /dev/vx/dmp/c2t6d0s2
```

Based on your knowledge of how the serial split brain condition came about, you must choose one disk's configuration to be used to import the disk group. For example, the following command imports the disk group using the configuration copy that is on side 0 of the split:

```
# /usr/sbin/vxdg -o selectcp=1045852127.32.olancha import newdg
```

When you have selected a preferred configuration copy, and the disk group has been imported, VxVM resets the serial IDs to 0 for the imported disks. The actual and expected serial IDs for any disks in the disk group that are not imported at this time remain unaltered.

Disabling a disk group

To disable a disk group, unmount and stop any volumes in the disk group, and then use the following command to deport it:

```
# vxdg deport diskgroup
```

Deporting a disk group does not actually remove the disk group. It disables use of the disk group by the system. Disks in a deported disk group can be reused, reinitialized, added to other disk groups, or imported for use on other systems. Use the `vxdg import` command to re-enable access to the disk group.

Destroying a disk group

The `vxdg` command provides a `destroy` option that removes a disk group from the system and frees the disks in that disk group for reinitialization:

```
# vxdg destroy diskgroup
```

Warning: This command destroys all data on the disks.

When a disk group is destroyed, the disks that are released can be re-used in other disk groups.

Recovering a destroyed disk group

If a disk group has been accidentally destroyed, you can recover it, provided that the disks that were in the disk group have not been modified or reused elsewhere.

To recover a destroyed disk group

- 1 Enter the following command to find out the disk group ID (`dgid`) of one of the disks that was in the disk group:

```
# vxdisk -s list disk_access_name
```

The disk must be specified by its disk access name, such as `c0t12d0s2`.

Examine the output from the command for a line similar to the following that specifies the disk group ID.

```
dgid:    963504895.1075.bass
```

- 2 Use the disk group ID to import the disk group:

```
# vxdg import dgid
```

Backing up and restoring disk group configuration data

The disk group configuration backup and restoration feature allows you to back up and restore all configuration data for disk groups, and for VxVM objects such as volumes that are configured within the disk groups. The `vxconfigbackupd` daemon monitors changes to the VxVM configuration and automatically records any configuration changes that occur. By default, `vxconfigbackup` stores 5 copies of the configuration backup and restoration (cbr) data. You can customize the number of cbr copies, between 1 to 5 copies.

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

VxVM provides the utilities, `vxconfigbackup` and `vxconfigrestore`, for backing up and restoring a VxVM configuration for a disk group.

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

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

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

Working with existing ISP disk groups

The Intelligent Storage Provisioning (ISP) feature of Veritas Volume Manager (VxVM) has been deprecated. This release does not support creating ISP disk groups. If you have existing ISP disk groups, you can import the disk groups without upgrading the disk group version. In this case, you cannot perform any operations on ISP volumes that would result in a configuration change. In addition, you cannot use any of the current release functionality that requires the upgraded disk group version.

You can upgrade an ISP disk group to the current disk group version. This operation converts all ISP volumes to standard (non-ISP) volumes and deletes ISP-specific objects. The ISP-specific objects include st pool, volume template, capability, and rules. This operation does not affect non-ISP volumes.

Note: When you upgrade the ISP disk group, all intent and storage pools information is lost. Only upgrade the disk group when this condition is acceptable.

To determine whether a disk group is an ISP disk group

- ◆ Check for the presence of storage pools, using the following command:

```
# vxprint
```

Sample output:

```
Disk group: mydg
TY NAME      ASSOC      KSTATE  LENGTH  PLOFFS  STATE      TUTILO  PUTILO
dg mydg      mydg      -       -       -       ALLOC_SUP  -       -

dm mydg2     ams_wms0_359 -       4120320 -       -       -       -
dm mydg3     ams_wms0_360 -       4120320 -       -       -       -

st mypool    -         -       -       -       DATA      -       -
dm mydg1     ams_wms0_358 -       4120320 -       -       -       -

v  myvol0    fsgen     ENABLED  20480   -       ACTIVE     -       -
pl myvol0-01 myvol0    ENABLED  20480   -       ACTIVE     -       -
sd mydg1-01  myvol0-01 ENABLED  20480   0       -         -       -

v  myvol1    fsgen     ENABLED  20480   -       ACTIVE     -       -
pl myvol1-01 myvol1    ENABLED  20480   -       ACTIVE     -       -
sd mydg1-02 myvol1-01 ENABLED  20480   0       -         -       -
```

In the sample output, `st mypool` indicates that `mydg` is an ISP disk group.

To upgrade an ISP disk group

- ◆ Upgrade the ISP disk group using the following command:

```
# vxdg upgrade ISP_diskgroup
```

To use an ISP disk group as is

- ◆ To import an ISP disk group, use the following command:

```
# vxdg import ISP_diskgroup
```

The ISP volumes in the disk group are not allowed to make any configuration changes until the disk group is upgraded. Attempting any operations such as grow shrink, add mirror, disk group split join, etc, on ISP volumes would give the following error:

This disk group is a ISP disk group. Dg needs to be migrated to non-ISP dg to allow any configuration changes. Please upgrade the dg to perform the migration.

Note: Non-ISP or VxVM volumes in the ISP disk group are not affected.

Operations that still work on ISP disk group without upgrading:

- Setting, removing, and replacing volume tags.
- Renaming of any VxVM objects such as volume, dg, plex, etc.
- Plex attach and detach.
- The `vxconfigbackup` and `vxconfigrestore` command can be used at the cost of losing any intent information

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 661.

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 659.

- 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 660.

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 disable 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.

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 vol01-02
```

In this case, the state of `vol01-02` 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 663.

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 667.

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 673.

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 671.

- 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 554.

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 673.

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 671.

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 681.

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:

```
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 679.

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 675.

Once encapsulated, the root disk can also be mirrored by using the `vxdiskadm` command.

See [“Mirroring an encapsulated root disk”](#) on page 684.

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 Solaris x64 systems, mirroring a root disk creates a GRUB boot menu entry for the Primary and Alternate (mirror) Boot disk.

For Solaris 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 686.

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 Solaris 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 687.

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 690.

`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 685.

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 685.

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.

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 696.

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 694.

The quota soft limit can be exceeded when VxFS preallocates space to a file.

See [“About extent attributes”](#) on page 179.

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 696.

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 `vxedquota(1M)`, `vxrepquota(1M)`, `vxquot(1M)`, `vxquota(1M)`, `vxquotaon(1M)`, and `vxquotaoff(1M)` 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 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 704.

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.

Reference

- [Chapter 34. Reverse path name lookup](#)
- [Appendix A. Tunable parameters](#)
- [Appendix B. Veritas File System disk layout](#)
- [Appendix C. Command reference](#)

Reverse path name lookup

This chapter includes the following topics:

- [About reverse path name lookup](#)

About 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.

Tunable parameters

This appendix includes the following topics:

- [About tuning Veritas Storage Foundation](#)
- [Tuning the VxFS file system](#)
- [DMP tunable parameters](#)
- [Methods to change Veritas Dynamic Multi-Pathing tunable parameters](#)
- [Tunable parameters for VxVM](#)
- [Methods to change Veritas Volume Manager tunable parameters](#)

About tuning Veritas Storage Foundation

Veritas Storage Foundation (SF) is widely used in a range of environments where performance plays a critical role. SF has a number of tunable parameters and configuration options that are meant to enable customization of the stack for the particular environment and workload in which SF is used. This guide helps administrators understand how some of these options affect performance, and provides guidelines for tuning the options.

Tuning the VxFS file system

This section describes the following kernel tunable parameters in VxFS:

- [Tuning inode table size](#)
- [Tuning performance optimization of inode allocation](#)
- [vx_maxlink](#)
- [Veritas Volume Manager maximum I/O size](#)

- [Partitioned directories](#)

Tuning inode table size

Tuning the internal inode table size includes the following tasks:

- Increasing the internal inode table size
- Changing the size of the directory name lookup cache

VxFS caches inodes in an inode table. The tunable for VxFS to determine the number of entries in its inode table is `vxfs_ninode`.

VxFS uses the value of `vxfs_ninode` in `/etc/system` as the number of entries in the VxFS inode table. By default, the file system uses a value of `vxfs_ninode`, which is computed based on system memory size.

To increase the internal inode table size

- 1 Open the `/etc/system` file.
- 2 Update the following line in the `/etc/system` file:

```
set vxfs:vxfs_ninode = new_value
```

Tuning performance optimization of inode allocation

The `delicache_enable` tunable parameter specifies whether performance optimization of inode allocation and inode reuse during a new file creation is turned on or off. The `delicache_enable` tunable is not supported for cluster file systems. You can specify the following values for `delicache_enable`:

- 0 – Disables delicache optimization.
- 1 – Enables delicache optimization.

The default value of `delicache_enable` is 1 for local mounts and 0 for cluster file systems.

Partitioned directories

You can enable or disable the partitioned directories feature by setting the `pdir_enable` tunable. Specifying a value of 1 enables partitioned directories, while specifying a value of 0 disables partitioned directories. The default value is 1.

You can set the `pdir_threshold` tunable to specify the threshold value in terms of directory size in bytes beyond which VxFS will partition a directory if you enabled partitioned directories. The default value is 32768.

The `-d` option of the `fsadm` command removes empty hidden directories from partitioned directories. If you disabled partitioned directories, the `fsadm -d` command also converts partitioned directories to regular directories.

The partitioned directories feature operates only on disk layout Version 8 or later file systems.

Warning: If the directories are huge, conversion between partitioned directories and regular directories or vice versa needs some time. If you enable the feature when the root directory already contains a large number of files, the conversion can occur at file system mount time, and can cause the mount to take a long time. Symantec recommends that the conversion is performed when directories are slightly or not populated.

vx_maxlink

The VxFS `vx_maxlink` tunable determines the number of sub-directories that can be created under a directory.

A VxFS file system obtains the value of `vx_maxlink` from the system configuration file `/etc/system`. By default, `vx_maxlink` is 32K. To change the computed value of `vx_maxlink`, you can add an entry to the system configuration file. For example:

```
set vxfs:vx_maxlink = 65534
```

sets `vx_maxlink` to the maximum number of sub-directories. Valid values are 1 to 65534 (FFFE hexadecimal). Changes to `vx_maxlink` take effect after rebooting.

Veritas Volume Manager maximum I/O size

When using VxFS with Veritas Volume Manager (VxVM), VxVM by default breaks up I/O requests larger than 256K. When using striping, to optimize performance, the file system issues I/O requests that are up to a full stripe in size. If the stripe size is larger than 256K, those requests are broken up.

To avoid undesirable I/O breakup, you can increase the maximum I/O size by changing the value of the `vol_maxio` parameter in the `/etc/system` file.

See [“vol_maxio”](#) on page 712.

vol_maxio

The `vol_maxio` parameter controls the maximum size of logical I/O operations that can be performed without breaking up a request. Logical I/O requests larger than this value are broken up and performed synchronously. Physical I/Os are broken

up based on the capabilities of the disk device and are unaffected by changes to the `vol_maxio` logical request limit.

Raising the `vol_maxio` limit can cause problems if the size of an I/O requires more memory or kernel mapping space than exists. The recommended maximum for `vol_maxio` is 20% of the smaller of physical memory or kernel virtual memory. It is not advisable to go over this limit. Within this limit, you can generally obtain the best results by setting `vol_maxio` to the size of your largest stripe. This applies to both RAID-0 striping and RAID-5 striping.

To increase the value of `vol_maxio`, add an entry to `/etc/system` (after the entry `forceload:drv/vxio`) and reboot for the change to take effect. For example, the following line sets the maximum I/O size to 16 MB:

```
set vxio:vol_maxio=32768
```

This parameter is in 512-byte sectors and is stored as a 16-bit number, so it cannot be larger than 65535.

DMP tunable parameters

DMP provides various parameters that you can use to tune your environment.

[Table A-1](#) shows the DMP parameters that can be tuned. You can set a tunable parameter online, without a reboot.

Table A-1 DMP parameters that are tunable

Parameter	Description
<code>dmp_cache_open</code>	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>.

Table A-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_daemon_count</code>	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.
<code>dmp_delayq_interval</code>	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.
<code>dmp_fast_recovery</code>	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>.
<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.

Table A-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<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 A-1 DMP parameters that are tunable (*continued*)

Parameter	Description
<code>dmp_lun_retry_timeout</code>	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.
<code>dmp_monitor_fabric</code>	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 A-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 IOs 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 A-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 Dynamic Multi-Pathing is installed as a component of Storage Foundation, the default value is <code>off</code>. When Veritas Dynamic Multi-Pathing is installed as a stand-alone product, 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 of 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 I/O policy” on page 234.

Table A-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. See “Configuring DMP path restoration policies” on page 247.
<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>vxddmpadm start restore</code> command. See “Configuring DMP path restoration policies” on page 247.

Table A-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>vxdmpadm start restore</code> command. See “Configuring DMP path restoration policies” on page 247.
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 247. 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 249.
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 A-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.

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 713.

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 721.
- Use the template method of the `vxddmpadm` command. See [“About tuning Veritas Dynamic Multi-Pathing \(DMP\) with templates”](#) on page 722.

Changing the values of DMP parameters with the vxddmpadm settune command line

To set a DMP timable 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 722.

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. DMP provides 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_CLARiion
    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.

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.

```
# vxddmpadm 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 713.

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)

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 730.

- **cvm**
Parameters to tune Cluster Volume Manager (CVM).
See “[Tunable parameters for CVM](#)” on page 742.
- **fmr**
Parameters to tune the FlashSnap functionality (FMR).
See “[Tunable parameters for FlashSnap \(FMR\)](#)” on page 736.
- **vvr**
Parameters to tune Veritas Volume Replicator (VVR).
See “[Tunable parameters for VVR](#)” on page 742.

Tunable parameters for core VxVM

[Table A-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 A-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.

Table A-2 Kernel tunable parameters for core VxVM (*continued*)

Parameter	Description
<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.
<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>voliomem_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.

Table A-2 Kernel tunable parameters for core VxVM (*continued*)

Parameter	Description
<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>voliomem_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).
<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.

Table A-2 Kernel tunable parameters for core VxVM (*continued*)

Parameter	Description
<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.
<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.

Table A-2 Kernel tunable parameters for core VxVM (*continued*)

Parameter	Description
<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 A-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 A-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 A-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 A-3 Kernel tunable parameters for FlashSnap (FMR)

Parameter	Description
<code>vol_fmr_logsz</code>	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 A-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 A-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 A-3 Kernel tunable parameters for FlashSnap (FMR) *(continued)*

Parameter	Description
<code>volpagemod_max_memsz</code>	

Table A-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 A-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 A-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 A-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 83.

Tunable parameters for VVR

[Table A-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 A-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 A-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 4MB (4194304 bytes). 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.

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 729.

Change the VxVM tunable parameters with one of the following methods:

Use the `vxtune` 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 `vxtune` command line”](#) on page 745.

Use the template method of the `vxtune` command.

See [“Changing the value of the Veritas Volume Manager tunable parameters using templates”](#) on page 747.

Use the `vxio.conf` file.

See [“Changing the tunable values using the `vxio.conf` file”](#) on page 748.

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.

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 Volume Replicator
vol_cmpres_threads	10	10	N	Maximum number of compression threads for VERITAS Volume Replicator
vol_default_iodelay	50	50	Y	Time to pause between I/O requests from VxVM utilities (10ms units)
vol_fmr_logsz	4	4	Y	Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes)
vol_max_adminio_poolsz	67108864	67108864	Y	Maximum amount of memory used by VxVM admin IO's (bytes)
.				
.				
.				

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.

- 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
```

- 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 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.

Veritas File System disk layout

This appendix includes the following topics:

- [About Veritas File System disk layouts](#)
- [VxFS Version 7 disk layout](#)
- [VxFS Version 8 disk layout](#)
- [VxFS Version 9 disk layout](#)

About Veritas File System 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 6	Version 6 disk layout enables features such as multi-volume support, cross-platform data sharing, named data streams, and File Change Log.	Deprecated
	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.	
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 694.

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 45.

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)

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 45.

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 45.

Command reference

This appendix includes the following topics:

- [Command completion for Veritas commands](#)
- [Veritas Volume Manager command reference](#)
- [Veritas Volume Manager manual pages](#)
- [Veritas File System command summary](#)
- [Veritas File System manual pages](#)

Command completion for Veritas commands

Veritas Storage Foundation supports command completion for Veritas Volume Manager (VxVM) commands and Dynamic Multi-Pathing (DMP) commands.

In this release, command completion is supported only on the bash shell. The shell must be bash version 2.4 or later.

To use this feature, press **Tab** while entering a supported VxVM or DMP command. The command is completed as far as possible. When there is a choice, the command completion displays the next valid options for the command. Enter one of the displayed values. A value in brackets indicates a user-specified value.

Note: Platform-specific options are not supported with command completion.

By default, you can use the command completion feature by invoking the bash shell on every log in. If you want to permanently enable the command completion, use the following command:

```
# vxctl cmdcompletion enable
```

The enable command completion creates the `.bash_profile` file, if it is not present.

To permanently disable the command completion, use the following command:

```
# vxctl cmdcompletion disable
```

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

The following commands support command completion:

- vxassist
- vxdisk
- vxplex
- vxprint
- vxsnap
- vxstat
- vxtune
- vxcache
- vxconfigd
- vxtask
- vxreattach
- vxdumpadm
- vxddladm
- vxvol
- vxcdsconvert
- vxresize
- vxctl
- vxsd
- vxdisksetup
- vxdiskunsetup
- vxrecover
- vxedit
- vxdg
- vxclustadm

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 777.

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 C-1](#) lists commands for obtaining information about objects in VxVM.
- [Table C-2](#) lists commands for administering disks.
- [Table C-3](#) lists commands for creating and administering disk groups.
- [Table C-4](#) lists commands for creating and administering subdisks.
- [Table C-5](#) lists commands for creating and administering plexes.
- [Table C-6](#) lists commands for creating volumes.
- [Table C-7](#) lists commands for administering volumes.
- [Table C-8](#) lists commands for monitoring and controlling tasks in VxVM.

Table C-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 262. Example: <pre># vxdisk -g mydg list</pre>
<code>vx dg list [<i>diskgroup</i>]</code>	Lists information about disk groups. See “Displaying disk group information” on page 623. Example: <pre># vx dg list mydg</pre>
<code>vx dg -s list</code>	Lists information about shared disk groups. Example: <pre># vx dg -s list</pre>
<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>vxinfo [-g <i>diskgroup</i>] [<i>volume</i> ...]</code>	Displays information about the accessibility and usability of volumes. See the <i>Veritas Storage Foundation and High Availability Troubleshooting Guide</i>. Example: <pre># vxinfo -g mydg myvol1 \ myvol2</pre>

Table C-1 Obtaining information about objects in VxVM (*continued*)

Command	Description
<code>vxprint -hrt [-g <i>diskgroup</i>] [<i>object</i> ...]</code>	Prints single-line information about objects in VxVM. Example: # vxprint -g mydg myvol1 \ myvol2
<code>vxlist</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.
<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 C-2 Administering disks

Command	Description
<code>vxdisk [-o full] reclaim {<i>disk enclosure 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 282. Example: # vxdiskadd c0t1d0

Table C-2 Administering disks (*continued*)

Command	Description
<code>vxedit [-g <i>diskgroup</i>] rename \ olddisk newdisk</code>	Renames a disk under control of VxVM. See “Renaming a disk” on page 285. Example: <pre># vxedit -g mydg rename \ mydg03 mydg02</pre>
<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 553. See “Making a disk available for hot-relocation use” on page 554. Examples: <pre># vxedit -g mydg set \ nohotuse=on mydg03 # vxedit -g mydg set \ nohotuse=off mydg03</pre>

Table C-2 Administering disks (*continued*)

Command	Description
<code>vxedit [-g <i>diskgroup</i>] set \</code> <code>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 551. See “Removing a disk from use as a hot-relocation spare” on page 552. 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>
<code>vxdbg -g <i>diskgroup</i> rmdisk <i>diskname</i></code>	Removes a disk from its disk group. See “Removing a disk from a disk group” on page 626. Example: <pre># vxdbg -g mydg rmdisk mydg02</pre>
<code>vxdiskunsetup <i>devicename</i></code>	Removes a disk from control of VxVM. See “Removing a disk from a disk group” on page 626. Example: <pre># vxdiskunsetup c0t3d0</pre>

Table C-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 disk group” on page 625. Example: # <code>vxdg init mydg \</code> <code>mydg01=c0t1d0</code>
<code>vxdg -g <i>diskgroup</i> listssbinfo</code>	Reports conflicting configuration information. See “Handling conflicting configuration copies” on page 648. Example: # <code>vxdg -g mydg listssbinfo</code>
<code>vxdg [-n <i>newname</i>] deport <i>diskgroup</i></code>	Deports a disk group and optionally renames it. See “Deporting a disk group” on page 627. Example: # <code>vxdg -n newdg deport mydg</code>
<code>vxdg [-n <i>newname</i>] import <i>diskgroup</i></code>	Imports a disk group and optionally renames it. See “Importing a disk group” on page 628. Example: # <code>vxdg -n newdg import mydg</code>
<code>vxdg [-n <i>newname</i>] -s import <i>diskgroup</i></code>	Imports a disk group as shared by a cluster, and optionally renames it. Example: # <code>vxdg -n newsdg -s import \</code> <code>mysdg</code>

Table C-3 Creating and administering disk groups (*continued*)

Command	Description
<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 593. 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 595. 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 598. Example: <pre># vxdg -o expand split mydg \ newdg myvol2 myvol3</pre>
<code>vxdg join <i>sourcedg</i> <i>targetdg</i></code>	Joins two disk groups. See “Joining disk groups” on page 599. Example: <pre># vxdg join newdg mydg</pre>
<code>vxdg -g <i>diskgroup</i> set \ <i>activation=ew ro sr sw off</i></code>	Sets the activation mode of a shared disk group in a cluster. Example: <pre># vxdg -g mysdg set \ activation=sw</pre>

Table C-3 Creating and administering disk groups (*continued*)

Command	Description
<code>vxrecover -g diskgroup -sb</code>	Starts all volumes in an imported disk group. See “Moving disk groups between systems” on page 631. Example: <pre># vxrecover -g mydg -sb</pre>
<code>vx dg destroy diskgroup</code>	Destroys a disk group and releases its disks. See “Destroying a disk group” on page 655. Example: <pre># vx dg destroy mydg</pre>

Table C-4 Creating and administering subdisks

Command	Description
<code>vxmake [-g diskgroup] sd subdisk \ diskname,offset,length</code>	Creates a subdisk. Example: <pre># vxmake -g mydg sd \ mydg02-01 mydg02,0,8000</pre>
<code>vxsd [-g diskgroup] assoc plex \ subdisk...</code>	Associates subdisks with an existing plex. Example: <pre># vxsd -g mydg assoc home-1 \ mydg02-01 mydg02-00 \ mydg02-01</pre>

Table C-4 Creating and administering subdisks (*continued*)

Command	Description
<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: <pre># vxsd -g mydg assoc \ vol101-01 mydg10-01:0 \ mydg11-01:1 mydg12-01:2</pre>
<code>vxsd [-g diskgroup] mv oldsubdisk \ newsubdisk ...</code>	Replaces a subdisk. Example: <pre># vxsd -g mydg mv mydg01-01 \ mydg02-01</pre>
<code>vxsd [-g diskgroup] -s size split \ subdisk sd1 sd2</code>	Splits a subdisk in two. Example: <pre># vxsd -g mydg -s 1000m \ split mydg03-02 mydg03-02 \ mydg03-03</pre>
<code>vxsd [-g diskgroup] join \ sd1 sd2 ... subdisk</code>	Joins two or more subdisks. Example: <pre># vxsd -g mydg join \ mydg03-02 mydg03-03 \ mydg03-02</pre>
<code>vxassist [-g diskgroup] move \ volume \!olddisk newdisk</code>	Relocates subdisks in a volume between disks. Example: <pre># vxassist -g mydg move \ myvol \!mydg02 mydg05</pre> Note: The ! character is a special character in some shells. This example shows how to escape it in a bash shell.

Table C-4 Creating and administering subdisks (*continued*)

Command	Description
<code>vxunreloc [-g <i>diskgroup</i>] <i>original_disk</i></code>	Relocates subdisks to their original disks. See “Moving relocated subdisks using vxunreloc” on page 555. Example: # vxunreloc -g mydg mydg01
<code>vxsd [-g <i>diskgroup</i>] dis <i>subdisk</i></code>	Dissociates a subdisk from a plex. Example: # vxsd -g mydg dis mydg02-01
<code>vxedit [-g <i>diskgroup</i>] rm <i>subdisk</i></code>	Removes a subdisk. Example: # vxedit -g mydg rm mydg02-01
<code>vxsd [-g <i>diskgroup</i>] -o rm dis <i>subdisk</i></code>	Dissociates and removes a subdisk from a plex. Example: # vxsd -g mydg -o rm dis \ mydg02-01

Table C-5 Creating and administering plexes

Command	Description
<code>vxmake [-g <i>diskgroup</i>] plex <i>plex</i> \ sd=<i>subdisk1</i>[,<i>subdisk2</i>,...]</code>	Creates a concatenated plex. Example: # vxmake -g mydg plex \ vol01-02 \ sd=mydg02-01,mydg02-02

Table C-5 Creating and administering plexes (*continued*)

Command	Description
<code>vxmake [-g diskgroup] plex <i>plex</i> \</code> <code>layout=stripe raid5 stwidth=<i>W</i> \</code> <code>ncolumn=<i>N</i> \</code> <code>sd=subdisk1[,subdisk2,...]</code>	Creates a striped or RAID-5 plex. Example: # <code>vxmake -g mydg plex pl-01 \</code> <code>layout=stripe stwidth=32 \</code> <code>ncolumn=2 \</code> <code>sd=mydg01-01,mydg02-01</code>
<code>vxplex [-g diskgroup] att volume <i>plex</i></code>	Attaches a plex to an existing volume. See “Reattaching a plex manually” on page 660. Example: # <code>vxplex -g mydg att vol101 \</code> <code>vol101-02</code>
<code>vxplex [-g diskgroup] det <i>plex</i></code>	Detaches a plex. Example: # <code>vxplex -g mydg det vol101-02</code>
<code>vxmend [-g diskgroup] off <i>plex</i></code>	Takes a plex offline for maintenance. Example: # <code>vxmend -g mydg off vol102-02</code>
<code>vxmend [-g diskgroup] on <i>plex</i></code>	Re-enables a plex for use. See “Reattaching a plex manually” on page 660. Example: # <code>vxmend -g mydg on vol102-02</code>
<code>vxplex [-g diskgroup] mv <i>oldplex</i> \</code> <code><i>newplex</i></code>	Replaces a plex. Example: # <code>vxplex -g mydg mv \</code> <code>vol102-02 vol102-03</code>

Table C-5 Creating and administering plexes (*continued*)

Command	Description
<code>vxplex [-g <i>diskgroup</i>] cp <i>volume</i> \ <i>newplex</i></code>	Copies a volume onto a plex. Example: <pre># vxplex -g mydg cp vol102 \ vol103-01</pre>
<code>vxmend [-g <i>diskgroup</i>] fix clean <i>plex</i></code>	Sets the state of a plex in an unstartable volume to CLEAN. See “Reattaching a plex manually” on page 660. Example: <pre># vxmend -g mydg fix clean \ vol102-02</pre>
<code>vxplex [-g <i>diskgroup</i>] -o rm dis <i>plex</i></code>	Dissociates and removes a plex from a volume. Example: <pre># vxplex -g mydg -o rm dis \ vol103-01</pre>

Table C-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: <pre># vxassist -g mydg maxsize \ layout=raid5 nlog=2</pre>

Table C-6 Creating volumes (*continued*)

Command	Description
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume length [layout=<i>layout</i>] \ [<i>attributes</i>]</code>	Creates a volume. See “Creating a volume on specific disks” on page 148. Example: # <code>vxassist -b -g mydg make \ myvol 20g layout=concat \ mydg01 mydg02</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume length layout=mirror \ [nmirror=<i>N</i>] [<i>attributes</i>]</code>	Creates a mirrored volume. See “Creating a mirrored volume” on page 143. Example: # <code>vxassist -b -g mydg make \ mymvol 20g layout=mirror \ nmirror=2</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume length 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. Example: # <code>vxassist -b -g mysdg make \ mysmvol 20g layout=mirror \ exclusive=on</code>
<code>vxassist -b [-g <i>diskgroup</i>] make \ volume length layout={stripe raid5} \ [stripeunit=<i>W</i>] [ncol=<i>N</i>] \ [<i>attributes</i>]</code>	Creates a striped or RAID-5 volume. See “Creating a striped volume” on page 145. See “Creating a RAID-5 volume” on page 146. Example: # <code>vxassist -b -g mydg make \ mysvol 20g layout=stripe \ stripeunit=32 ncol=4</code>

Table C-6 Creating volumes (*continued*)

Command	Description
<code>vxassist -b [-g <i>diskgroup</i>] make \</code> <code>volume length layout=mirror \</code> <code>mirror=ctlr [<i>attributes</i>]</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 C-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 611. Example: # <code>vxassist -g mydg mirror \</code> # <code>myvol mydg10</code>

Table C-7 Administering volumes (*continued*)

Command	Description
<code>vxassist [-g <i>diskgroup</i>] remove \ mirror <i>volume</i> [<i>attributes</i>]</code>	Removes a mirror from a volume. See “Removing a mirror” on page 615. Example: # vxassist -g mydg remove \ mirror myvol \!mydg11 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>] \ {growto growby} <i>volume length</i></code>	Grows a volume to a specified size or by a specified amount. Example: # vxassist -g mydg growby \ myvol 10g
<code>vxassist [-g <i>diskgroup</i>] \ {shrinkto shrinkby} <i>volume length</i></code>	Shrinks a volume to a specified size or by a specified amount. Example: # vxassist -g mydg shrinkto \ myvol 20g
<code>vxresize -b -F vxfs [-g <i>diskgroup</i>] \ <i>volume length diskname ...</i></code>	Resizes a volume and the underlying Veritas File System. Example: # vxresize -b -F vxfs \ -g mydg myvol 20g mydg10 \ mydg11

Table C-7 Administering volumes (*continued*)

Command	Description
<code>vxsnap [-g <i>diskgroup</i>] prepare <i>volume</i> [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 343. Example: <pre># vxsnap -g mydg prepare \ myvol drl=on</pre>
<code>vxsnap [-g <i>diskgroup</i>] make \ source=<i>volume</i>\ /newvol=<i>snapvol</i>\ [/nmirror=<i>number</i>]</code>	Takes a full-sized instant snapshot of a volume by breaking off plexes of the original volume. See “Creating instant snapshots” on page 342. Example: <pre># vxsnap -g mydg make \ source=myvol/\ newvol=mysnpvol/\ nmirror=2</pre>
<code>vxsnap [-g <i>diskgroup</i>] make \ source=<i>volume</i>/snapvol=<i>snapvol</i></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 347. See “Creating instant snapshots” on page 342. Example: <pre># vxsnap -g mydg make \ source=myvol/snapvol=snpvol</pre>

Table C-7 Administering volumes (*continued*)

Command	Description
<pre>vxmake [-g diskgroup] cache \ cache_object cachevolname=volume \ [regionsize=size]</pre>	Creates a cache object for use by space-optimized instant snapshots. See “Creating a shared cache object” on page 345. 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>
<pre>vxsnap [-g diskgroup] make \ source=volume/newvol=snapvol\ /cache=cache_object</pre>	Takes a space-optimized instant snapshot of a volume. See “Creating instant snapshots” on page 342. Example: <pre># vxsnap -g mydg make \ source=myvol/\ newvol=mysosvol/\ cache=cobj</pre>
<pre>vxsnap [-g diskgroup] refresh snapshot</pre>	Refreshes a snapshot from its original volume. See “Refreshing an instant space-optimized snapshot” on page 364. Example: <pre># vxsnap -g mydg refresh \ mysnpvol</pre>

Table C-7 Administering volumes (*continued*)

Command	Description
<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 366. Example: <pre># vxsnap -g mydg dis mysnpvol</pre>
<code>vxsnap [-g <i>diskgroup</i>] unprepare \ volume</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 \ volume [layout=<i>layout</i>] \ [relayout_<i>options</i>]</code>	Performs online relayout of a volume. See “Performing online relayout” on page 605. Example: <pre># vxassist -g mydg relayout \ vol2 layout=stripe</pre>
<code>vxassist [-g <i>diskgroup</i>] relayout \ volume layout=raid5 \ stripeunit=<i>W</i> \ ncol=<i>N</i></code>	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 605. Example: <pre># vxassist -g mydg relayout \ vol3 layout=raid5 \ stripeunit=16 ncol=4</pre>

Table C-7 Administering volumes (*continued*)

Command	Description
<code>vxrelayout [-g <i>diskgroup</i>] -o bg \</code> <code>reverse <i>volume</i></code>	Reverses the direction of a paused volume relayout. See “Volume sets” on page 97. Example: # vxrelayout -g mydg -o bg \ reverse vol3
<code>vxassist [-g <i>diskgroup</i>] convert \</code> <code><i>volume</i> [layout=<i>layout</i>] \</code> <code>[<i>convert_options</i>]</code>	Converts between a layered volume and a non-layered volume layout. Example: # vxassist -g mydg convert \ vol3 layout=stripe-mirror
<code>vxassist [-g <i>diskgroup</i>] remove \</code> <code><i>volume volume</i></code>	Removes a volume. See “Removing a volume” on page 662. Example: # vxassist -g mydg remove \ myvol

Table C-8 Monitoring and controlling tasks

Command	Description
<code>command [-g <i>diskgroup</i>] -t <i>tasktag</i> \</code> <code>[<i>options</i>] [<i>arguments</i>]</code>	Specifies a task tag to a VxVM command. See “Specifying task tags” on page 601. Example: # vxrecover -g mydg \ -t mytask -b mydg05

Table C-8 Monitoring and controlling tasks (*continued*)

Command	Description
<code>vxtask [-h] [-g <i>diskgroup</i>] list</code>	Lists tasks running on a system. See “Using the vxtask command” on page 603. Example: <pre># vxtask -h -g mydg list</pre>
<code>vxtask monitor <i>task</i></code>	Monitors the progress of a task. See “Using the vxtask command” on page 603. Example: <pre># vxtask monitor mytask</pre>
<code>vxtask pause <i>task</i></code>	Suspends operation of a task. See “Using the vxtask command” on page 603. Example: <pre># vxtask pause mytask</pre>
<code>vxtask -p [-g <i>diskgroup</i>] list</code>	Lists all paused tasks. See “Using the vxtask command” on page 603. Example: <pre># vxtask -p -g mydg list</pre>
<code>vxtask resume <i>task</i></code>	Resumes a paused task. See “Using the vxtask command” on page 603. Example: <pre># vxtask resume mytask</pre>

Table C-8 Monitoring and controlling tasks (*continued*)

Command	Description
<code>vxtask abort task</code>	Cancels a task and attempts to reverse its effects. See “Using the vxtask command” on page 603. Example: <pre># vxtask abort mytask</pre>

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 C-9](#) lists the manual pages in section 1M for commands that are used to administer Veritas Volume Manager.

Table C-9 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.

Table C-9 Section 1M manual pages (*continued*)

Name	Description
<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>vxdc</code>	Perform operations on version 0 DCO objects and DCO volumes.
<code>vxctl</code>	Control the volume configuration daemon.
<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.

Table C-9 Section 1M manual pages (*continued*)

Name	Description
<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>vxdlmpadm</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.
<code>vxiod</code>	Start, stop, and report on Veritas Volume Manager kernel I/O threads.
<code>vxlufinish</code>	SolarisFinish 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.

Table C-9 Section 1M manual pages (*continued*)

Name	Description
<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.
<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.

Table C-9 Section 1M manual pages (*continued*)

Name	Description
<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 C-10](#) lists the manual pages in section 4 that describe the format of files that are used by Veritas Volume Manager.

Table C-10 Section 4 manual pages

Name	Description
<code>vol_pattern</code>	Disk group search specifications.
<code>vxmake</code>	<code>vxmake</code> description file.

Section 7 — device driver interfaces

[Table C-11](#) lists the manual pages in section 7 that describe the interfaces to Veritas Volume Manager devices.

Table C-11 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 C-12](#) describes the VxFS-specific commands.

Table C-12 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>fcladm</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
<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.

Table C-12 VxFS commands (*continued*)

Command	Description
<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.
<code>glmdump</code>	Reports stuck Group Lock Managers (GLM) locks in a cluster file system.
<code>glmstat</code>	Group Lock Managers (GLM) statistics gathering utility. This functionality is available only with the Veritas Storage Foundation Cluster File System High Availability product.
<code>ls</code>	Lists files and directories on a VxFS file system.
<code>mkdstfs</code>	SmartTier file system creation utility.
<code>mkfs</code>	Constructs a VxFS file system.
<code>mount</code>	Mounts a VxFS file system.
<code>mv</code>	Moves files and directories on a VxFS file system.
<code>ncheck</code>	Generates path names from inode numbers for a VxFS file system.
<code>qioadmin</code>	Administers VxFS Quick I/O for Databases cache.
<code>qiomkfile</code>	Creates a VxFS Quick I/O device file. This functionality is available only with the Veritas Quick I/O for Databases feature.
<code>qiostat</code>	Displays statistics for VxFS Quick I/O for Databases. This functionality is available only with the Veritas Quick I/O for Databases feature.
<code>setext</code>	Sets extent attributes on a file in a VxFS file system.
<code>umount</code>	Unmounts a VxFS file system.

Table C-12 VxFS commands (*continued*)

Command	Description
<code>vxcompress</code>	Compresses and uncompresses files.
<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>vxfilesnap</code>	Makes a copy-on-write copy of a file in a VxFS file system.
<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.

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 C-13](#) describes the VxFS-specific section 1 manual pages.

Table C-13 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>fsmmap</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.
<code>qiomkfile</code>	Creates a VxFS Quick I/O device file. This functionality is available only with the Veritas Quick I/O for Databases feature.
<code>qiostat</code>	Displays statistics for VxFS Quick I/O for Databases. This functionality is available only with the Veritas Quick I/O for Databases feature.
<code>setext</code>	Sets extent attributes on a file in a VxFS file system.
<code>vxcompress</code>	Compresses or uncompresses files.
<code>vxfilesnap</code>	Makes a copy-on-write copy of a file in a VxFS file system.

[Table C-14](#) describes the VxFS-specific section 1M manual pages.

Table C-14 Section 1M manual pages

Section 1M	Description
<code>df_vxfs</code>	Reports the number of free disk blocks and inodes for a VxFS file system.
<code>fcladm</code>	Administers VxFS File Change Logs.
<code>ff_vxfs</code>	Lists file names and inode information for a VxFS file system.
<code>fsadm_vxfs</code>	Resizes or reorganizes a VxFS file system.
<code>fsapadm</code>	Administers VxFS allocation policies.
<code>fscat_vxfs</code>	Cats a VxFS file system.

Table C-14 Section 1M manual pages (*continued*)

Section 1M	Description
<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_vxfs</code>	Checks and repairs a VxFS file system.
<code>fsckptadm</code>	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.
<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 Cluster File System product.
<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>fsppadm</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.

Table C-14 Section 1M manual pages (*continued*)

Section 1M	Description
<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.
<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 C-15](#) describes the VxFS-specific section 3 manual pages.

Table C-15 Section 3 manual pages

Section 3	Description
<code>vxfs_ap_alloc2</code>	Allocates an <code>fsap_info2</code> structure.

Table C-15 Section 3 manual pages (*continued*)

Section 3	Description
<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.
<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.

Table C-15 Section 3 manual pages (*continued*)

Section 3	Description
<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.
<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_ioffsets</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.

Table C-15 Section 3 manual pages (*continued*)

Section 3	Description
<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 C-16](#) describes the VxFS-specific section 4 manual pages.

Table C-16 Section 4 manual pages

Section 4	Description
<code>fs_vxfs</code>	Provides the format of a VxFS file system volume.
<code>inode_vxfs</code>	Provides the format of a VxFS file system inode.
<code>tunefstab</code>	Describes the VxFS file system tuning parameters table.

[Table C-17](#) describes the VxFS-specific section 7 manual pages.

Table C-17 Section 7 manual pages

Section 7	Description
<code>vxfsio</code>	Describes the VxFS file system control functions.

Index

Symbols

/dev/vx/dmp directory 102
/dev/vx/rdmp directory 102
/etc/default/vxassist file 120, 554
/etc/default/vxdg file 626
/etc/vfstab file 662
/etc/volboot file 59
/etc/vx/darecs file 59
/etc/vx/dmppolicy.info file 234
/etc/vx/volboot file 633
/kernel/drv/vxio.conf file 748
/kernel/drv/vxspec.conf file 748
/lib/svc/method/vxvm-recover file 558

A

A/A disk arrays 100
A/A-A disk arrays 101
A/P disk arrays 101
A/P-C disk arrays 101–102
A/P-F disk arrays 101
A/P-G disk arrays 102
A5x00 arrays
 removing and replacing disks 673
access control lists 35
access port 101
active path attribute 231
active paths
 devices 232–233
ACTIVE state 336
Active/Active disk arrays 100
Active/Passive disk arrays 101
adaptive load-balancing 234
adaptiveminq policy 234
adding disks 282
allocation policies 181
 default 180
 extent 29
 extent based 28, 32
APM
 configuring 249

array policy module (APM)
 configuring 249
array ports
 disabling for DMP 240
 displaying information about 220
 enabling for DMP 241
array support library (ASL) 190
Array Volume ID
 device naming 265
arrays
 DMP support 189
ASL
 array support library 189–190
Asymmetric Active/Active disk arrays 101
ATTACHING state 336
attributes
 active 231
 autogrow 345, 350
 autogrowby 345
 cache 350
 cachesize 350
 dcolen 95, 384
 for specifying storage 148
 highwatermark 345
 maxautogrow 345
 maxdev 636
 mirdg 358
 mirvol 358
 ncachemirror 350
 ndcomirror 384
 ndcomirs 343
 newvol 356
 nmirror 356
 nomanual 231
 nopreferred 231
 preferred priority 231
 primary 231
 secondary 231
 setting for paths 230, 233
 snapvol 352, 358
 source 352, 358
 standby 232

attributes *(continued)*
 syncing 342, 370

autogrow
 tuning 372
 autogrow attribute 345, 350
 autogrowby attribute 345
 autotrespass mode 101

B

backups
 created using snapshots 342
 creating for volumes 315
 creating using instant snapshots 342
 creating using third-mirror snapshots 374
 for multiple volumes 359, 379
 of disk group configuration 656
 bad block revectoring 163
 balanced path policy 235
 base minor number 634
 blkclear mount option 163
 block based architecture 41
 blockmap for a snapshot file system 417
 blocks on disks 55
 boot command
 -a flag 687
 syntax 687
 boot disk
 encapsulating 684
 listing volumes on 688
 mirroring 684
 unencapsulating 692
 boot disks
 alternate 685
 boot-time restrictions 682
 booting root volumes 682
 BROKEN state 336
 buffered file systems 33
 buffered I/O 295

C

c# 49
 cache
 for space-optimized instant snapshots 317
 cache advisories 297
 cache attribute 350
 cache objects
 creating 345
 enabling 346

cache objects *(continued)*
 listing snapshots in 372
 caches
 creating 345
 deleting 374
 finding out snapshots configured on 374
 growing 373
 listing snapshots in 372
 removing 374
 resizing 373
 shrinking 373
 stopping 374
 used by space-optimized instant snapshots 317
 cachesize attribute 350
 campus clusters
 serial split brain condition in 648
 cascade instant snapshots 337
 cascaded snapshot hierarchies
 creating 363
 categories
 disks 190
 CDS
 compatible disk groups 626
 cds attribute 626
 check_all policy 247
 check_alternate policy 248
 check_disabled policy 248
 check_periodic policy 248
 checkpoint interval 730
 cio
 Concurrent I/O 169
 clone_disk flag 638
 cloned disks 637–638
 cluster mount 37
 clusters
 use of DMP in 110
 vol_fmr_logsz tunable 737
 columns
 changing number of 608
 in striping 64
 mirroring in striped-mirror volumes 146
 CommandCentral Storage 97
 commands
 cron 45
 fsadm 44
 getext 183
 setext 183
 compressing files 41
 concatenated volumes 61, 142

- concatenated-mirror volumes
 - creating 144
 - defined 69
 - recovery 143
- concatenation 61
- configuration backup and restoration 656
- configuration changes
 - monitoring using vxnotify 604
- configuration database
 - listing disks with 639
 - metadata 639
 - reducing size of 588
- Configuring DMP
 - using templates 722
- contiguous reservation 182
- Controller ID
 - displaying 219
- controllers
 - disabling for DMP 240
 - disabling in DMP 208
 - displaying information about 218
 - enabling for DMP 241
 - mirroring across 152
 - number 49
 - specifying to vxassist 148
- converting a data Storage Checkpoint to a nodata Storage Checkpoint 393
- convosync mount option 160, 165
- copy-on-write
 - used by instant snapshots 335
- copy-on-write technique 322, 388
- copymaps 93–94
- cp_vxfs 183
- cpio_vxfs 183
- creating a multi-volume support file system 462
- creating file systems with large files 167
- creating files with mkfs 155, 158
- cron 44, 177
- cron sample script 178
- customized naming
 - DMP nodes 266

D

- d# 49
- data change object
 - DCO 93
- data copy 294
- data redundancy 67, 70
- data Storage Checkpoints definition 325
- data synchronous I/O 164, 295
- data transfer 294
- data volume configuration 84
- database replay logs and sequential DRL 83
- databases
 - integrity of data in 315
 - resilvering 83
 - resynchronizing 83
- DCO
 - adding version 0 DCOs to volumes 383
 - considerations for disk layout 594
 - data change object 93
 - dissociating version 0 DCOs from volumes 386
 - effect on disk group split and join 594
 - instant snap version 93
 - log plexes 91
 - log volume 93
 - moving log plexes 385
 - reattaching version 0 DCOs to volumes 386
 - removing version 0 DCOs from volumes 386
 - specifying storage for version 0 plexes 385
 - used with DRL 83
 - version 0 93
 - version 20 93–94
 - versioning 93
- dcolon attribute 95, 384
- DDI_NT_FABRIC property 188
- DDL 103
 - Device Discovery Layer 193
- default
 - allocation policy 180
- defaulttdg 585
- defragmentation 44
 - extent 177
 - scheduling with cron 177
- delaylog 34
- delaylog mount option 161
- device discovery
 - introduced 103
 - partial 188
- Device Discovery Layer 193
- Device Discovery Layer (DDL) 103, 193
- device names 49
 - configuring persistent 267
 - user-specified 266
- device nodes
 - controlling access for volume sets 457
 - displaying access for volume sets 457
 - enabling access for volume sets 456

- device nodes *(continued)*
 - for volume sets 455
- devices
 - adding foreign 204
 - fabric 188
 - JBOD 189
 - listing all 194
 - making invisible to VxVM 205
 - nopriv 680
 - path redundancy 232–233
- direct data transfer 294
- direct I/O 294
- directory reorganization 178
- dirty flags set on volumes 81
- dirty region logging.. *See* DRL
- dirty regions 738
- disabled file system
 - snapshot 331
- disabled paths 210
- discovered direct I/O 295
- disk access records
 - stored in /etc/vx/darecs 59
- disk arrays
 - A/A 100
 - A/A-A 101
 - A/P 101
 - A/P-F 101
 - A/P-G 102
 - Active/Active 100
 - Active/Passive 101
 - adding disks to DISKS category 201
 - Asymmetric Active/Active 101
 - defined 50
 - excluding support for 199
 - JBOD devices 189
 - listing excluded 200
 - listing supported 198
 - listing supported disks in DISKS category 200
 - multipathed 59
 - re-including support for 200
 - removing disks from DISKS category 203
 - supported with DMP 198
- Disk Group Split/Join 318
- disk groups
 - avoiding conflicting minor numbers on import 634
 - clearing locks on disks 632
 - compatible with CDS 626
 - configuration backup and restoration 656
 - creating with old version number 623
 - disk groups *(continued)*
 - defined 53
 - deporting 627
 - destroying 655
 - determining the default disk group 584
 - disabling 655
 - displaying boot disk group 585
 - displaying default disk group 585
 - displaying free space in 624
 - displaying information about 623
 - displaying version of 622
 - features supported by version 618
 - forcing import of 633
 - free space in 549
 - importing 628
 - importing with cloned disks 638
 - ISP 657
 - joining 590, 599
 - layout of DCO plexes 594
 - limitations of move
 - split. *See* and join
 - listing objects affected by a move 593
 - moving between systems 631
 - moving disks between 587, 596
 - moving licensed EMC disks between 596
 - moving objects between 588, 595
 - recovering destroyed 655
 - recovery from failed reconfiguration 592
 - removing disks from 626
 - renaming 646
 - reorganizing 588
 - reserving minor numbers 634
 - restarting moved volumes 600
 - root 54
 - rootdg 54
 - serial split brain condition 648
 - setting default disk group 585
 - splitting 589, 598
 - upgrading version of 622
 - version 618, 622
 - disk layout
 - Version 6 751
 - Version 7 751
 - Version 8 751
 - Version 9 751
 - disk media names 54
 - disk names
 - configuring persistent 267
 - disk## 55

disk### ## 55

disks 190

- adding 282
- adding to DISKS category 201
- array support library 190
- categories 190
- changing naming scheme 264
- clearing locks on 632
- cloned 638
- complete failure messages 548
- configuring newly added 187
- configuring persistent names 267
- determining failed 548
- Device Discovery Layer 193
- disabled path 210
- discovery of by DMP 187
- discovery of by VxVM 189
- disk access records file 59
- disk arrays 50
- displaying information 262–263
- displaying information about 262, 624
- displaying naming scheme 265
- displaying spare 551
- dynamic LUN expansion 258
- enabled path 210
- encapsulation 675, 681
- enclosures 104
- excluding free space from hot-relocation use 553
- failure handled by hot-relocation 545
- formatting 272
- handling clones 637
- handling duplicated identifiers 637
- hot-relocation 543
- initializing 273
- installing 272
- invoking discovery of 191
- layout of DCO plexes 594
- listing tags on 639
- listing those supported in JBODs 200
- making available for hot-relocation 551
- making free space available for hot-relocation use 554
- marking as spare 551
- mirroring boot disk 684
- mirroring root disk 684
- mirroring volumes on 612
- moving between disk groups 587, 596
- moving disk groups between systems 631
- moving volumes from 586

disks (*continued*)

- nopriv devices 680
- number 49
- OTHER_DISKS category 190
- partial failure messages 547
- postponing replacement 667
- primary path 210
- reinitializing 281
- releasing from disk groups 655
- removing 282, 667
- removing from A5x00 arrays 673
- removing from disk groups 626
- removing from DISKS category 203
- removing from pool of hot-relocation spares 552
- removing from VxVM control 626, 663
- removing tags from 640
- removing with subdisks 284–285
- renaming 285
- replacing 667
- replacing in A5x00 arrays 673
- replacing removed 671
- root disk 681
- scanning for 188
- secondary path 210
- setting tags on 638
- spare 549
- specifying to vxassist 148
- UDID flag 637
- unique identifier 637
- VM 54
- writing a new identifier to 638
- DISKS category 190
 - adding disks 201
 - listing supported disks 200
 - removing disks 203
- displaying
 - DMP nodes 214
 - HBA information 219
 - redundancy levels 232
 - supported disk arrays 198
- displaying mounted file systems 175
- displaying statistics
 - erroneous I/Os 228
 - queued I/Os 228
- DMP
 - booting from DMP devices 110
 - check_all restore policy 247
 - check_alternate restore policy 248
 - check_disabled restore policy 248

DMP (continued)

- check_periodic restore policy 248
- configuring disk devices 187
- configuring DMP path restoration policies 247
- configuring I/O throttling 244
- configuring response to I/O errors 243, 246
- disabling array ports 240
- disabling controllers 240
- disabling multi-pathing 205
- disabling paths 240
- disk discovery 187
- displaying DMP database information 208
- displaying DMP node for a path 213
- displaying DMP node for an enclosure 213–214
- displaying DMP nodes 214
- displaying information about array ports 220
- displaying information about controllers 218
- displaying information about enclosures 219
- displaying information about paths 209
- displaying LUN group for a node 215
- displaying paths controlled by DMP node 216
- displaying paths for a controller 216
- displaying paths for an array port 217
- displaying recoveryoption values 246
- displaying status of DMP path restoration
 - thread 249
- displaying TPD information 221
- dynamic multi-pathing 100
- enabling array ports 241
- enabling controllers 241
- enabling multi-pathing 207
- enabling paths 241
- enclosure-based naming 103
- gathering I/O statistics 225
- in a clustered environment 110
- load balancing 108
- logging levels 715
- metanodes 102
- nodes 102
- path aging 714
- path failover mechanism 107
- path-switch tunable 718
- renaming an enclosure 242
- restore policy 247
- scheduling I/O on secondary paths 237
- setting the DMP restore polling interval 247
- stopping the DMP restore daemon 249
- tuning with templates 722
- vxdmpadm 211

DMP nodes

- displaying consolidated information 214
- setting names 266

DMP support

- JBOD devices 189
- dmp_cache_open tunable 713
- dmp_daemon_count tunable 714
- dmp_delayq_interval tunable 714
- dmp_fast_recovery tunable 714
- dmp_health_time tunable 714
- dmp_log_level tunable 715
- dmp_low_impact_probe 715
- dmp_lun_retry_timeout tunable 716
- dmp_monitor_fabric tunable 716
- dmp_monitor_osevent tunable 717
- dmp_monitor_ownership tunable 717
- dmp_native_multipathing tunable 717
- dmp_native_support tunable 718
- dmp_path_age tunable 718
- dmp_pathswitch_blks_shift tunable 718
- dmp_probe_idle_lun tunable 719
- dmp_probe_threshold tunable 719
- dmp_restore_cycles tunable 719
- dmp_restore_interval tunable 719
- dmp_restore_state tunable 720
- dmp_scsi_timeout tunable 720
- dmp_sfg_threshold tunable 721
- dmp_stat_interval tunable 721

DR

- dynamic reconfiguration 109

DRL

- dirty region logging 82
- hot-relocation limitations 545
- log subdisks 83
- maximum number of dirty regions 738
- minimum number of sectors 738
- sequential 83
- use of DCO with 83

- dynamic LUN expansion 258

- dynamic reconfiguration 109

E**EMC arrays**

- moving disks between disk groups 596

EMC PowerPath

- coexistence with DMP 192

EMC Symmetrix

- autodiscovery 192

- enabled paths
 - displaying 210
- encapsulating disks 675, 681
- encapsulating volumes 460
- encapsulation
 - failure of 679
- enclosure-based naming 104, 264
 - DMP 103
- enclosures 104
 - displaying information about 219
 - issues with nopriv disks 269
 - issues with simple disks 269
 - path redundancy 232–233
 - setting attributes of paths 230, 233
- enhanced data integrity modes 33
- ENOSPC 402
- erroneous I/Os
 - displaying statistics 228
- error messages
 - Disk for disk group not found 633
 - Disk group has no valid configuration copies 633
 - Disk group version doesn't support feature 617
 - Disk is in use by another host 632
 - Disk is used by one or more subdisks 626
 - Disk not moving
 - but subdisks on it are 593
 - import failed 632
 - It is not possible to encapsulate 679
 - No valid disk found containing disk group 632
 - The encapsulation operation failed 679
 - tmpsize too small to perform this relayout 77
 - unsupported layout 679
 - vxvg listmove failed 593
- errord daemon 106
- errors
 - handling transient errors 716
- expansion 45
- explicit failover mode 101
- extent 28, 32–33, 179
 - attributes 33, 179
 - reorganization 178
- extent allocation 29
 - aligned 180
 - control 33, 179
 - fixed size 180
- extent attributes 33, 179
- external quotas file 694

F

- fabric devices 188
- FAILFAST flag 107
- failover mode 101
- failure handled by hot-relocation 545
- failure in RAID-5 handled by hot-relocation 545
- FastResync
 - effect of growing volume on 95
 - limitations 96
 - Non-Persistent 89
 - Persistent 89–90, 315
 - size of bitmap 737
 - snapshot enhancements 333
 - use with snapshots 88
- FastResync/cache object metadata cache size
 - tunable 741
- fc_woff 703
- file
 - sparse 181
- file compression 41
- file system
 - block size 185
 - buffering 33
 - displaying mounted 175
- file systems
 - mirroring on root disk 687
 - unmounting 662
- fileset
 - primary 320
- FileSnaps
 - about 327
 - data mining, reporting, and testing 412
 - virtual desktops 411
 - write intensive applications 412
 - backup 329
 - best practices 411
 - block map fragmentation 329
 - concurrent I/O 328
 - copy-on-write 328
 - creation 409
 - properties 327
 - reading from 329
 - using 410
- fixed extent size 180
- fixed write size 181
- FlashSnap 313
- FMR.. See FastResync
- foreign devices
 - adding 204

- formatting disks 272
- fragmentation
 - monitoring 177–178
 - reorganization facilities 177
 - reporting 177
- fragmented file system characteristics 177
- free space in disk groups 549
- free space monitoring 176
- freeze 297
- freezing and thawing, relation to Storage Checkpoints 320
- fsadm 44
 - how to minimize file system free space fragmentation 173
 - how to reorganize a file system 173
 - how to resize a file system 171
 - reporting extent fragmentation 178
 - scheduling defragmentation using cron 178
 - thin reclamation 437
- fsadm_vxfs 169
- fsck 416
- fsck 393
- fsckptadm
 - Storage Checkpoint administration 389
- fstyp
 - how to determine the file system type 175
- fsvoladm 462
- full-sized instant snapshots 334
 - creating 352
 - creating volumes for use as 347
- fullinst snapshot type 369

G

- get I/O parameter ioctl 298
- getext 183
- getfacl 35

H

- HBA information
 - displaying 219
- HBAs
 - listing ports 195
 - listing supported 194
 - listing targets 195
- highwatermark attribute 345
- hot-relocation
 - complete failure messages 548
 - configuration summary 550

- hot-relocation (*continued*)
 - daemon 544
 - defined 82
 - detecting disk failure 545
 - detecting plex failure 545
 - detecting RAID-5 subdisk failure 545
 - excluding free space on disks from use by 553
 - limitations 545
 - making free space on disks available for use by 554
 - marking disks as spare 551
 - modifying behavior of 558
 - notifying users other than root 559
 - operation of 543
 - partial failure messages 547
 - preventing from running 559
 - reducing performance impact of recovery 559
 - removing disks from spare pool 552
 - subdisk relocation 550
 - subdisk relocation messages 555
 - unrelocating subdisks 555
 - unrelocating subdisks using vxunreloc 556
 - use of free space in disk groups 549
 - use of spare disks 549
 - use of spare disks and free space 549
 - using only spare disks for 554
 - vxrelocd 544
- how to access a Storage Checkpoint 391
- how to create a Storage Checkpoint 390
- how to determine the file system type 175
- how to display mounted file systems 170
- how to minimize file system free space fragmentation 173
- how to mount a Storage Checkpoint 391
- how to remove a Storage Checkpoint 391
- how to reorganize a file system 173
- how to resize a file system 171
- how to unmount a Storage Checkpoint 393

I

- I/O
 - direct 294
 - gathering statistics for DMP 225
 - kernel threads 48
 - scheduling on secondary paths 237
 - sequential 295
 - synchronous 295
 - throttling 107

- I/O operations
 - maximum size of 732
- I/O policy
 - displaying 233
 - example 238
 - specifying 234
- I/O requests
 - asynchronous 164
 - synchronous 163
- I/O throttling 244
- I/O throttling options
 - configuring 247
- identifiers for tasks 601
- idle LUNs 719
- implicit failover mode 101
- Importing
 - ISP disk group 657
- initialization
 - of disks 273
- inode table 711
 - internal 711
 - sizes 711
- inodes, block based 29
- instant snap version
 - of DCOs 93
- instant snapshots
 - backing up multiple volumes 359
 - cascaded 337
 - creating backups 342
 - creating for volume sets 360
 - creating full-sized 352
 - creating space-optimized 349
 - creating volumes for use as full-sized 347
 - displaying information about 368
 - dissociating 366
 - full-sized 334
 - improving performance of synchronization 371
 - reattaching 364
 - refreshing 364
 - removing 367
 - restoring volumes using 366
 - space-optimized 317
 - splitting hierarchies 367
 - synchronizing 370
- intent log 27, 33
 - multi-volume file system support 460
- intent log resizing 28
- intent logging 316
- internal inode table 711

- internal quotas file 694
- ioctl calls 732–733
- iSCSI parameters
 - administering with DDL 197
 - setting with vxddladm 197
- ISP
 - disk groups 657
- ISP disk group
 - Importing 657
 - Upgrading 657

J

- JBOD
 - DMP support 189
- JBODs
 - adding disks to DISKS category 201
 - listing supported disks 200
 - removing disks from DISKS category 203

K

- kernel tunable parameters 710

L

- large files 34, 167
 - creating file systems with 167
 - mounting file systems with 168
- largefiles mount option 168
- layered volumes
 - defined 74, 143
 - striped-mirror 68
- layouts
 - left-symmetric 72
 - types of volume 142
- left-symmetric layout 72
- link objects 336
- linked break-off snapshots 336
 - creating 357
- linked third-mirror snapshots
 - reattaching 365
- listing
 - alternate boot disks 685
 - DMP nodes 214
 - supported disk arrays 198
- load balancing 100
 - displaying policy for 233
 - specifying policy for 234
- local mount 37
- lock clearing on disks 632

- log mount option 159, 161
- log subdisks
 - DRL 83
- logdisk 147
- logical units 101
- logiosize mount option 163
- logs
 - RAID-5 74, 81
 - specifying number for RAID-5 147
- LUN 101
- LUN expansion 258
- LUN group failover 102
- LUN groups
 - displaying details of 215
- LUNs
 - idle 719
 - thin provisioning 427

M

- maxautogrow attribute 345
- maxdev attribute 636
- maximum I/O size 712
- memory
 - granularity of allocation by VxVM 734
 - maximum size of pool for VxVM 734
 - minimum size of pool for VxVM 736
 - persistence of FastResync in 89
- messages
 - complete disk failure 548
 - hot-relocation of subdisks 555
 - partial disk failure 547
- metadata 639
 - multi-volume support 460
- metanodes
 - DMP 102
- migrating to thin storage 427
- mincache mount option 160, 163
- minimum queue load balancing policy 236
- minimum redundancy levels
 - displaying for a device 232
 - specifying for a device 233
- minor numbers 634
- mirbrk snapshot type 369
- mirdg attribute 358
- mirrored volumes
 - changing read policies for 153
 - configuring VxVM to create by default 612
 - creating 143
 - creating across controllers 152

- mirrored volumes *(continued)*
 - creating across targets 150
 - defined 142
 - dirty region logging 82
 - DRL 82
 - FastResync 82
 - FR 82
 - logging 82
 - snapshots 88
- mirrored-concatenated volumes
 - creating 144
 - defined 67
- mirrored-stripe volumes
 - benefits of 67
 - creating 145
 - defined 143
- mirroring
 - boot disk 684
 - defined 67
 - root disk 684
- mirroring plus striping 68
- mirrors
 - adding to volumes 611
 - creating snapshot 376
 - defined 58
 - removing from volumes 615
 - specifying number of 144
- mirvol attribute 358
- mirvol snapshot type 369
- mkfs
 - creating files with 155, 158
 - creating large files 169
- modes
 - enhanced data integrity 33
- monitoring fragmentation 177
- mount 169
 - how to display mounted file systems 170
 - mounting a Storage Checkpoint 391
 - pseudo device 392
- mount options 157
 - blkclear 163
 - choosing 157
 - combining 169
 - convosync 160, 165
 - delaylog 34, 161
 - extended 33
 - largefiles 168
 - log 159, 161
 - logiosize 163

- mount options (*continued*)
 - mincache 160, 163
 - nodatainlog 159, 163
 - tmplog 162
- mounted file system
 - displaying 175
- mounting a file system
 - option combinations 169
 - with large files 168
- mounting a Storage Checkpoint 393
- mounting a Storage Checkpoint of a cluster file system 393
- mrl
 - keyword 232
- multi-pathing
 - disabling 205
 - displaying information about 209
 - enabling 207
- multi-volume file systems 459
- Multi-Volume Support 451
- multi-volume support
 - creating a MVS file system 462
- multiple block operations 29
- mv_vxfs 183

N

- name space
 - preserved by Storage Checkpoints 389
- names
 - changing for disk groups 646
 - defining for snapshot volumes 379
 - device 49
 - disk media 54
 - plex 57
 - renaming disks 285
 - subdisk 55
 - VM disk 55
 - volume 57
- naming
 - DMP nodes 266
- naming scheme
 - changing for disks 264
 - changing for TPD enclosures 268
 - displaying for disks 265
- ncachemirror attribute 350
- ncheck 708
- ndcomirror attribute 384
- ndcomirs attribute 343
- newvol attribute 356
- nmirror attribute 355–356
- nodata Storage Checkpoints 393
- nodata Storage Checkpoints definition 325
- nodatainlog mount option 159, 163
- nodes
 - DMP 102
- nomannual path attribute 231
- non-autotrespass mode 101
- Non-Persistent FastResync 89
- nopreferred path attribute 231
- nopriv devices 680
- nopriv disks
 - issues with enclosures 269

O

- O_SYNC 160
- objects
 - physical 49
 - virtual 51
- online invalid status 263
- online relayout
 - changing number of columns 608
 - changing region size 611
 - changing speed of 611
 - changing stripe unit size 608
 - controlling progress of 610
 - defined 76
 - destination layouts 605
 - failure recovery 80
 - how it works 76
 - limitations 79
 - monitoring tasks for 610
 - pausing 610
 - performing 605
 - resuming 610
 - reversing direction of 611
 - specifying non-default 608
 - specifying plexes 609
 - specifying task tags for 609
 - temporary area 77
 - transformation characteristics 80
 - transformations and volume length 80
 - types of transformation 606
 - viewing status of 610
- online status 263
- OpenBoot PROMs (OPB) 687
- ordered allocation 147, 150
- OTHER_DISKS category 190

P

- parity in RAID-5 70
- partial device discovery 188
- partition size
 - displaying the value of 233
 - specifying 235
- partitions
 - number 49–50
 - slices 50
- path aging 714
- path failover in DMP 107
- pathgroups
 - creating 206
- paths
 - disabling for DMP 240
 - enabling for DMP 241
 - setting attributes of 230, 233
- performance
 - changing values of tunables 721, 744
 - displaying tunable values 749
 - improving for instant snapshot
 - synchronization 371
 - load balancing in DMP 108
 - overall 157
 - snapshot file systems 416
- persistence
 - device naming option 265
- persistent device name database 267
- persistent device naming 267
- Persistent FastResync 89–90, 95, 315
- physical disks
 - clearing locks on 632
 - complete failure messages 548
 - determining failed 548
 - displaying information 262
 - displaying information about 262, 624
 - displaying spare 551
 - excluding free space from hot-relocation use 553
 - failure handled by hot-relocation 545
 - installing 272
 - making available for hot-relocation 551
 - making free space available for hot-relocation
 - use 554
 - marking as spare 551
 - moving between disk groups 587, 596
 - moving disk groups between systems 631
 - moving volumes from 586
 - partial failure messages 547
 - postponing replacement 667
- physical disks *(continued)*
 - releasing from disk groups 655
 - removing 282, 667
 - removing from A5x00 arrays 673
 - removing from disk groups 626
 - removing from pool of hot-relocation spares 552
 - removing with subdisks 284–285
 - replacing 667
 - replacing in A5x00 arrays 673
 - replacing removed 671
 - spare 549
- physical objects 49
- ping-pong effect 110
- plex attribute 356
- plexes
 - adding to snapshots 381
 - changing read policies for 153
 - complete failure messages 548
 - converting to snapshot 378
 - defined 56
 - failure in hot-relocation 545
 - maximum number of subdisks 733
 - maximum number per volume 57
 - mirrors 58
 - moving 385
 - names 57
 - partial failure messages 547
 - putting online 660
 - reattaching 660
 - recovering after correctable hardware failure 548
 - removing from volumes 615
 - sparse 80
 - specifying for online relay layout 609
 - striped 63
 - types 56
- point-in-time copy solutions
 - applications 304
- polling interval for DMP restore 247
- ports
 - listing 195
- PowerPath
 - coexistence with DMP 192
- prefer read policy 154
- preferred plex
 - read policy 154
- preferred priority path attribute 231
- primary boot disk failure 685
- primary fileset relation to Storage Checkpoints 320
- primary path 101, 210

primary path attribute 231
 priority load balancing 236

PROMs

boot 687

prtconf

displaying tunables 749

pseudo device 392

Q

queued I/Os

displaying statistics 228

quota commands 695

quotacheck 696

quotas 693

exceeding the soft limit 694

hard limit 408, 693

soft limit 693

quotas file 694

quotas.grp file 694

R

RAID-0 63

RAID-0+1 67

RAID-1 67

RAID-1+0 68

RAID-5

hot-relocation limitations 545

logs 74, 81

parity 70

specifying number of logs 147

subdisk failure handled by hot-relocation 545

volumes 70

RAID-5 volumes

changing number of columns 608

changing stripe unit size 608

creating 147

defined 142

raw device nodes

controlling access for volume sets 457

displaying access for volume sets 457

enabling access for volume sets 456

for volume sets 455

read policies

changing 153

prefer 154

round 153

select 154

siteread 154

read policies (*continued*)

split 154

read-only Storage Checkpoints 391

recovery

checkpoint interval 730

I/O delay 730

recovery accelerator 83

recovery option values

configuring 246

redo log configuration 84

redundancy

of data on mirrors 142

of data on RAID-5 142

redundancy levels

displaying for a device 232

specifying for a device 233

redundant-loop access 105

regionsize attribute 343, 345

reinitialization of disks 281

relayout

changing number of columns 608

changing region size 611

changing speed of 611

changing stripe unit size 608

controlling progress of 610

limitations 79

monitoring tasks for 610

online 76

pausing 610

performing online 605

resuming 610

reversing direction of 611

specifying non-default 608

specifying plexes 609

specifying task tags for 609

storage 76

transformation characteristics 80

types of transformation 606

viewing status of 610

relocation

automatic 543

complete failure messages 548

limitations 545

partial failure messages 547

removable Storage Checkpoints definition 326

removing devices

from VxVM control 205

removing disks 667

removing physical disks 282

- reorganization
 - directory 178
 - extent 178
- replacing disks 667
- replay logs and sequential DRL 83
- report extent fragmentation 177
- reservation space 180
- resilvering
 - databases 83
- restoration of disk group configuration 656
- restore policy
 - check_all 247
 - check_alternate 248
 - check_disabled 248
 - check_periodic 248
- restored daemon 106
- resyncfromoriginal snapback 341
- resyncfromreplica snapback 341
- resynchronization
 - checkpoint interval 730
 - I/O delay 730
 - of volumes 81
- resynchronizing
 - databases 83
 - snapshots 318
- retry option values
 - configuring 246
- Reverse Path Name Lookup 708
- root disk
 - defined 681
 - encapsulating 684
 - listing volumes on 688
 - mirroring 684
 - mirroring other file systems on 687
 - unencapsulating 692
- root disk group 54
- root disks
 - booting alternate 685
- root volume 682
- rootability 681
 - removing 692
- rootdg 54
- round read policy 153
- round-robin
 - load balancing 236
 - read policy 153

S

s# 49–50

- SAN storage
 - configuring volumes on 97
- scandisks
 - vxdisk subcommand 188
- secondary path 101
- secondary path attribute 231
- secondary path display 210
- select read policy 154
- sequential DRL
 - defined 83
 - maximum number of dirty regions 738
- sequential I/O 295
- serial split brain condition
 - correcting 653
 - in campus clusters 648
 - in disk groups 648
- setext 183
- setfacl 35
- setting
 - path redundancy levels 233
- simple disks
 - issues with enclosures 269
- single active path policy 237
- siteread read policy 154
- slices
 - number 49
 - partitions 50
- SmartMove feature 427
- SmartSync 83
 - disabling on shared disk groups 742
 - enabling on shared disk groups 742
- SmartTier 451
 - multi-volume file system support 460
- snap objects 92
- snap volume naming 341
- snapabort 333
- snapback
 - defined 334
 - merging snapshot volumes 380
 - resyncfromoriginal 341
 - resyncfromreplica 341, 380
- snapclear
 - creating independent volumes 381
- snapiir snapshot type 369
- snapiof 419
- snapped file systems 36, 330
 - performance 416
 - unmounting 330
- snapiread 416

- snapshot file systems 36, 330
 - blockmap 417
 - creating 419
 - data block area 417
 - disabled 331
 - fscat 416
 - fuser 330
 - mounting 419
 - multiple 330
 - on cluster file systems 330
 - performance 416
 - read 416
 - super-block 417
- snapshot hierarchies
 - creating 363
 - splitting 367
- snapshot mirrors
 - adding to volumes 362
 - removing from volumes 363
- snapshots
 - adding mirrors to volumes 362
 - adding plexes to 381
 - and FastResync 88
 - backing up multiple volumes 359, 379
 - backing up volumes online using 342
 - cascaded 337
 - comparison of features 86
 - converting plexes to 378
 - creating a hierarchy of 363
 - creating backups using third-mirror 374
 - creating for volume sets 360
 - creating full-sized instant 352
 - creating independent volumes 381
 - creating instant 342
 - creating linked break-off 357
 - creating snapshots of 338
 - creating space-optimized instant 349
 - creating third-mirror break-off 354
 - creating volumes for use as full-sized instant 347
 - defining names for 379
 - displaying information about 382
 - displaying information about instant 368
 - dissociating instant 366
 - finding out those configured on a cache 374
 - full-sized instant 87, 334
 - hierarchy of 337
 - improving performance of synchronization 371
 - instant 317
 - linked break-off 336
 - snapshots *(continued)*
 - listing for a cache 372
 - merging with original volumes 380
 - of volumes 85
 - on multiple volumes 341
 - reattaching instant 364
 - reattaching linked third-mirror 365
 - refreshing instant 364
 - removing 378
 - removing instant 367
 - removing linked snapshots from volumes 363
 - removing mirrors from volumes 363
 - restoring from instant 366
 - resynchronization on snapback 341
 - resynchronizing 318
 - resynchronizing volumes from 380
 - space-optimized instant 317
 - synchronizing instant 370
 - third-mirror 86
 - use of copy-on-write mechanism 335
- snapsize 418
- snapstart 333
- snapvol attribute 352, 358
- snapwait 355, 358
- source attribute 352, 358
- space-optimized instant snapshots 317
 - creating 349
- spaceopt snapshot type 369
- spanned volumes 61
- spanning 61
- spare disks
 - displaying 551
 - marking disks as 551
 - used for hot-relocation 549
- sparse file 181
- sparse plexes 80
- specifying
 - redundancy levels 233
- split read policy 154
- standby path attribute 232
- states
 - of link objects 336
- statistics gathering 107
- storage
 - clearing 163
 - ordered allocation of 147, 150
 - uninitialized 163
- storage accounts 97
- storage attributes and volume layout 148

- storage cache 317
 - used by space-optimized instant snapshots 317
 - Storage Checkpoints 319
 - accessing 391
 - administration of 389
 - converting a data Storage Checkpoint to a nodata Storage Checkpoint with multiple Storage Checkpoints 396
 - creating 390
 - data Storage Checkpoints 325
 - definition of 388
 - difference between a data Storage Checkpoint and a nodata Storage Checkpoint 394
 - freezing and thawing a file system 320
 - mounting 391
 - multi-volume file system support 460
 - nodata Storage Checkpoints 325, 393
 - operation failures 402
 - pseudo device 392
 - read-only Storage Checkpoints 391
 - removable Storage Checkpoints 326
 - removing 391
 - space management 402
 - synchronous vs. asynchronous conversion 393
 - types of 324
 - unmounting 393
 - using the fsck command 393
 - writable Storage Checkpoints 391
 - storage groups 97
 - storage processor 101
 - storage relayout 76
 - stripe columns 64
 - stripe units
 - changing size 608
 - defined 64
 - stripe-mirror-col-split-trigger-pt 146
 - striped plexes
 - defined 63
 - striped volumes
 - changing number of columns 608
 - changing stripe unit size 608
 - creating 145
 - defined 142
 - failure of 63
 - specifying non-default number of columns 145
 - specifying non-default stripe unit size 145
 - striped-mirror volumes
 - benefits of 68
 - creating 146
 - striped-mirror volumes *(continued)*
 - defined 143
 - mirroring columns 146
 - mirroring subdisks 146
 - trigger point for mirroring 146
 - striping 63
 - striping plus mirroring 67
 - subdisk names 55
 - subdisks
 - blocks 55
 - complete failure messages 548
 - defined 55
 - determining failed 548
 - DRL log 83
 - hot-relocation 82, 543, 550
 - hot-relocation messages 555
 - listing original disks after hot-relocation 557
 - maximum number per plex 733
 - mirroring in striped-mirror volumes 146
 - moving after hot-relocation 555
 - partial failure messages 547
 - RAID-5 failure of 545
 - specifying different offsets for unrelocation 557
 - unrelocating after hot-relocation 555
 - unrelocating to different disks 556
 - unrelocating using vxunreloc 556
 - super-block 417
 - SVID requirement
 - VxFS conformance to 46
 - synchronization
 - controlling for instant snapshots 370
 - improving performance of 371
 - synchronous I/O 295
 - syncing attribute 342, 370
 - syncpause 370
 - syncresume 370
 - syncstart 370
 - syncstop 370
 - syncwait 370
 - system failure recovery 27, 33
 - system performance
 - overall 157
- ## T
- t# 49
 - tags
 - for tasks 601
 - listing for disks 639
 - removing from disks 640

- tags (*continued*)
 - removing from volumes 616
 - renaming 616
 - setting on disks 638
 - setting on volumes 616
 - specifying for online relay layout tasks 609
 - specifying for tasks 601
- target IDs
 - number 49
 - specifying to vxassist 148
- target mirroring 150
- targets
 - listing 195
- task monitor in VxVM 601
- tasks
 - aborting 603
 - changing state of 603
 - identifiers 601
 - listing 603
 - managing 602
 - modifying parameters of 603
 - monitoring 603
 - monitoring online relay layout 610
 - pausing 603
 - resuming 603
 - specifying tags 601
 - specifying tags on online relay layout operation 609
 - tags 601
- temporary area used by online relay layout 77
- temporary directories 34
- thaw 297
- thin provisioning
 - using 427
- Thin Reclamation 433
- thin reclamation
 - fsadm 437
- thin storage
 - using 427
- third-mirror
 - snapshots 86
- third-mirror break-off snapshots
 - creating 354
- third-party driver (TPD) 191
- throttling 107
- tmplog mount option 162
- TPD
 - displaying path information 221
 - support for coexistence 191
- tpdm mode attribute 268
- trigger point in striped-mirror volumes 146
- tunable I/O parameters
 - Volume Manager maximum I/O size 712
- tunables
 - changing values of 721, 744
 - displaying using prtconf 749
 - dmp_cache_open 713
 - dmp_daemon_count 714
 - dmp_delayq_interval 714
 - dmp_fast_recovery 714
 - dmp_health_time 714
 - dmp_log_level 715
 - dmp_low_impact_probe 715
 - dmp_lun_retry_timeout 716
 - dmp_monitor_fabric 716
 - dmp_monitor_osevent 717
 - dmp_monitor_ownership 717
 - dmp_native_multipathing 717
 - dmp_native_support 718
 - dmp_path_age 718
 - dmp_pathswitch_blks_shift 718
 - dmp_probe_idle_lun 719
 - dmp_probe_threshold 719
 - dmp_restore_cycles 719
 - dmp_restore_interval 719
 - dmp_restore_state 720
 - dmp_scsi_timeout 720
 - dmp_sfg_threshold 721
 - dmp_stat_interval 721
 - FastrResync/cache object metadata cache size 741
 - maximum vol_stats_enable 733
 - vol_checkpoint_default 730
 - vol_default_iodelay 730
 - vol_fmr_logsz 89, 737
 - vol_kmsg_resend_period 731
 - vol_kmsg_send_period 731
 - vol_max_vol 731
 - vol_maxio 732
 - vol_maxioctl 732
 - vol_maxparallelio 732
 - vol_maxspecialio 733
 - vol_subdisk_num 733
 - volcvm_smartsync 742
 - voldr_max_drtregs 738
 - voldr_max_seq_dirty 83, 738
 - voldr_min_regionsz 738
 - voliomem_chunk_size 734
 - voliomem_maxpool_sz 734

tunables (*continued*)

- voliot_errbuf_dflt 735
- voliot_iobuf_default 735
- voliot_iobuf_limit 735
- voliot_iobuf_max 735
- voliot_max_open 736
- volpagemod_max_memsz 741
- volraid_minpool_size 736
- volraid_rsrtransmax 736
- VVR

- changing. See using vxio.conf file

Tuning DMP

- using templates 722

tuning VxFS 710

U

- UDID flag 637
- udid_mismatch flag 637
- umount command 170
- unencapsulating the root disk 692
- uninitialized storage, clearing 163
- unmount 393
 - a snapped file system 330
- Upgrading
 - ISP disk group 657
- use_all_paths attribute 237
- use_avid
 - vxddladm option 265
- user-specified device names 266
- utility
 - vxtune 745, 747

V

- V-5-1-2829 617
- V-5-1-552 626
- V-5-1-587 632
- V-5-2-3091 593
- V-5-2-369 627
- V-5-2-4292 593
- Veritas Operations Manager 30
- version 0
 - of DCOs 93
- version 20
 - of DCOs 93–94
- Version 6 disk layout 751
- Version 7 disk layout 751
- Version 8 disk layout 751
- Version 9 disk layout 751

versioning

- of DCOs 93

versions

- disk group 622
- displaying for disk group 622
- upgrading 622

virtual disks 45

virtual objects 51

VM disks

- defined 54
- displaying spare 551
- excluding free space from hot-relocation use 553
- making free space available for hot-relocation use 554
- marking as spare 551
- mirroring volumes on 612
- moving volumes from 586
- names 55
- postponing replacement 667
- removing from pool of hot-relocation spares 552
- renaming 285

vol## 57

vol##-## 57

- vol_checkpoint_default tunable 730
- vol_default_iodelay tunable 730
- vol_fmr_logsz tunable 89, 737
- vol_kmsg_resend_period tunable 731
- vol_kmsg_send_period tunable 731
- vol_max_vol tunable 731
- vol_maxio tunable 732
- vol_maxio tunable I/O parameter 712
- vol_maxioctl tunable 732
- vol_maxparallelio tunable 732
- vol_maxspecialio tunable 733
- vol_subdisk_num tunable 733
- volbrk snapshot type 369
- volcvm_smartsync tunable 742
- voldrl_max_drtregs tunable 738
- voldrl_max_seq_dirty tunable 83, 738
- voldrl_min_regionsz tunable 738
- voliomem_chunk_size tunable 734
- voliomem_maxpool_sz tunable 734
- voliot_errbuf_dflt tunable 735
- voliot_iobuf_default tunable 735
- voliot_iobuf_limit tunable 735
- voliot_iobuf_max tunable 735
- voliot_max_open tunable 736
- volpagemod_max_memsz tunable 741
- volraid_minpool_size tunable 736

- volraid_rsrtransmax tunable 736
- volume resynchronization 81
- volume sets
 - adding volumes to 453
 - administering 451
 - controlling access to raw device nodes 457
 - creating 452
 - creating instant snapshots of 360
 - displaying access to raw device nodes 457
 - enabling access to raw device nodes 456
 - listing details of 453
 - raw device nodes 455
 - removing volumes from 453
 - starting 454
 - stopping 454
- volumes
 - adding mirrors 611
 - adding snapshot mirrors to 362
 - adding to volume sets 453
 - adding version 0 DCOs to 383
 - backing up 315
 - backing up online using snapshots 342
 - boot-time restrictions 682
 - booting root 682
 - changing layout online 605
 - changing number of columns 608
 - changing read policies for mirrored 153
 - changing stripe unit size 608
 - concatenated 61, 142
 - concatenated-mirror 69, 143
 - configuring on SAN storage 97
 - creating concatenated-mirror 144
 - creating for use as full-sized instant snapshots 347
 - creating from snapshots 381
 - creating mirrored 143
 - creating mirrored-concatenated 144
 - creating mirrored-stripe 145
 - creating RAID-5 147
 - creating snapshots 377
 - creating striped 145
 - creating striped-mirror 146
 - defined 57
 - displaying information about snapshots 382
 - dissociating version 0 DCOs from 386
 - effect of growing on FastResync maps 95
 - excluding storage from use by vxassist 148
 - flagged as dirty 81
 - layered 68, 74, 143
 - volumes *(continued)*
 - limit on number of plexes 57
 - limitations 57
 - listing on boot (root) disk 688
 - maximum number of 731
 - merging snapshots 380
 - mirrored 67, 142
 - mirrored-concatenated 67
 - mirrored-stripe 67, 143
 - mirroring across controllers 152
 - mirroring across targets 150
 - mirroring all 612
 - mirroring on disks 612
 - moving from VM disks 586
 - names 57
 - naming snap 341
 - performing online relayout 605
 - RAID-0 63
 - RAID-0+1 67
 - RAID-1 67
 - RAID-1+0 68
 - RAID-10 68
 - RAID-5 70, 142
 - reattaching plexes 660
 - reattaching version 0 DCOs to 386
 - recovering after correctable hardware failure 548
 - removing 662
 - removing from /etc/vfstab 662
 - removing linked snapshots from 363
 - removing mirrors from 615
 - removing plexes from 615
 - removing snapshot mirrors from 363
 - removing version 0 DCOs from 386
 - restarting moved 600
 - restoring from instant snapshots 366
 - resynchronizing from snapshots 380
 - snapshots 85
 - spanned 61
 - specifying non-default number of columns 145
 - specifying non-default relayout 608
 - specifying non-default stripe unit size 145
 - specifying storage for version 0 DCO plexes 385
 - specifying use of storage to vxassist 148
 - stopping activity on 662
 - striped 63, 142
 - striped-mirror 68, 143
 - taking multiple snapshots 341
 - trigger point for mirroring in striped-mirror 146
 - types of layout 142

- VX_DSYNC 295
- VX_FREEZE 297, 696
- VX_GETCACHE 297
- VX_SETCACHE 297
- VX_SNAPREAD 416
- VX_THAW 297
- VX_UNBUFFERED 295
- vxassist
 - adding DCOs to volumes 384
 - adding mirrors to volumes 611
 - creating cache volumes 345
 - creating concatenated-mirror volumes 144
 - creating mirrored volumes 144
 - creating mirrored-concatenated volumes 144
 - creating mirrored-stripe volumes 145
 - creating RAID-5 volumes 147
 - creating snapshots 375
 - creating striped volumes 145
 - creating striped-mirror volumes 146
 - creating volumes for use as full-sized instant snapshots 348
 - defaults file 120
 - defining layout on specified storage 148
 - displaying information about snapshots 382
 - dissociating snapshots from volumes 381
 - excluding storage from use 148
 - listing tags set on volumes 616
 - merging snapshots with volumes 380
 - mirroring across controllers 152
 - mirroring across targets 150, 152
 - mirroring file systems on root disk 687
 - moving DCO plexes 385
 - relaying out volumes online 605
 - removing mirrors 615
 - removing plexes 615
 - removing tags from volumes 616
 - removing version 0 DCOs from volumes 386
 - removing volumes 663
 - replacing tags set on volumes 616
 - resynchronizing volumes from snapshots 380
 - setting default values 120
 - setting tags on volumes 616–617
 - snapshot 333
 - snapback 334
 - snapshot 334
 - snapshot 333
 - specifying number of mirrors 144
 - specifying number of RAID-5 logs 147
 - specifying ordered allocation of storage 150
- vxassist (*continued*)
 - specifying plexes for online relayout 609
 - specifying storage attributes 148
 - specifying storage for version 0 DCO plexes 385
 - specifying tags for online relayout tasks 609
 - taking snapshots of multiple volumes 379
- vxcache
 - listing snapshots in a cache 372
 - resizing caches 373
 - starting cache objects 346
 - stopping a cache 374
 - tuning cache autogrow 373
- vxcached
 - tuning 372
- vxconfigd
 - managing with vxctl 58
 - monitoring configuration changes 604
- vxdaresore
 - handling simple/nopriv disk failures 269
- vxdc
 - dissociating version 0 DCOs from volumes 386
 - reattaching version 0 DCOs to volumes 386
 - removing version 0 DCOs from volumes 386
- vxctl
 - managing vxconfigd 58
 - setting default disk group 586
- vxctl enable
 - invoking device discovery 191
- vxddladm
 - adding disks to DISKS category 201
 - adding foreign devices 204
 - changing naming scheme 265
 - displaying the disk-naming scheme 265
 - listing all devices 194
 - listing configured devices 196–197
 - listing configured targets 195–196
 - listing excluded disk arrays 200–201
 - listing ports on a Host Bus Adapter 195
 - listing supported disk arrays 198
 - listing supported disks in DISKS category 200
 - listing supported HBAs 194
 - removing disks from DISKS category 192, 203–204
 - setting iSCSI parameters 197
 - used to exclude support for disk arrays 199
 - used to re-include support for disk arrays 200
- vxdg
 - clearing locks on disks 633

vx dg (continued)

- controlling CDS compatibility of new disk groups 626
- correcting serial split brain condition 654
- creating disk groups 625
- creating disk groups with old version number 623
- deporting disk groups 628
- destroying disk groups 655
- disabling a disk group 655
- displaying boot disk group 585
- displaying default disk group 585
- displaying disk group version 622
- displaying free space in disk groups 624
- displaying information about disk groups 623
- forcing import of disk groups 633
- importing a disk group containing cloned disks 638
- importing cloned disks 640
- importing disk groups 629
- joining disk groups 599
- listing disks with configuration database copies 639
- listing objects affected by move 593
- listing spare disks 551
- moving disk groups between systems 631
- moving disks between disk groups 587
- moving objects between disk groups 595
- placing a configuration database on cloned disks 639
- recovering destroyed disk groups 655
- removing disks from disk groups 626
- renaming disk groups 646
- setting base minor number 635
- setting maximum number of devices 636
- splitting disk groups 598
- upgrading disk group version 622

vx disk

- clearing locks on disks 633
- displaying information about disks 624
- displaying multi-pathing information 210
- listing disks 263
- listing spare disks 551
- listing tags on disks 639
- notifying dynamic LUN expansion 258
- placing a configuration database on a cloned disk 639
- removing tags from disks 640
- scanning disk devices 188
- setting tags on disks 638

vx disk (continued)

- updating the disk identifier 638

vx disk scandisks

- rescanning devices 188
- scanning devices 188

vx diskadd

- creating disk groups 625
- placing disks under VxVM control 282

vx diskadm

- Add or initialize one or more disks 273, 625
- adding disks 273
- changing the disk-naming scheme 264
- creating disk groups 625
- deporting disk groups 627
- Enable access to (import) a disk group 629
- Encapsulate one or more disks 677
- Exclude a disk from hot-relocation use 553
- excluding free space on disks from hot-relocation use 553
- importing disk groups 629
- initializing disks 273
- List disk information 263
- listing spare disks 551
- Make a disk available for hot-relocation use 554
- making free space on disks available for hot-relocation use 554
- Mark a disk as a spare for a disk group 552
- marking disks as spare 552
- Mirror volumes on a disk 612
- mirroring disks 682
- mirroring root disks 685
- mirroring volumes 612
- Move volumes from a disk 586
- moving disk groups between systems 634
- moving disks between disk groups 587
- moving subdisks from disks 627
- moving volumes from VM disks 586
- Remove a disk 283, 627
- Remove a disk for replacement 668
- Remove access to (deport) a disk group 627
- removing disks from pool of hot-relocation spares 553
- Replace a failed or removed disk 671
- Turn off the spare flag on a disk 553

vx diskconfig

- purpose of 187

vx diskunsetup

- removing disks from VxVM control 626, 663

- vxddmpadm
 - changing TPD naming scheme 268
 - configuring an APM 250
 - configuring I/O throttling 244
 - configuring response to I/O errors 243, 246
 - disabling controllers in DMP 208
 - disabling I/O in DMP 240
 - displaying APM information 250
 - displaying DMP database information 208
 - displaying DMP node for a path 213, 215
 - displaying DMP node for an enclosure 213–214
 - displaying I/O error recovery settings 246
 - displaying I/O policy 233
 - displaying I/O throttling settings 246
 - displaying information about controllers 218
 - displaying information about enclosures 219
 - displaying partition size 233
 - displaying paths controlled by DMP node 216
 - displaying status of DMP restoration thread 249
 - displaying TPD information 221
 - enabling I/O in DMP 241
 - gathering I/O statistics 225
 - listing information about array ports 220
 - removing an APM 250
 - renaming enclosures 242
 - setting I/O policy 236–237
 - setting path attributes 231
 - setting restore polling interval 247
 - specifying DMP path restoration policy 247
 - stopping DMP restore daemon 249
- vxddmpadm list
 - displaying DMP nodes 214
- vxddump 183
- vxedit
 - excluding free space on disks from hot-relocation use 553
 - making free space on disks available for hot-relocation use 554
 - marking disks as spare 551
 - removing a cache 374
 - removing disks from pool of hot-relocation spares 552
 - removing instant snapshots 367
 - removing snapshots from a cache 374
 - removing volumes 663
 - renaming disks 286
- VxFS
 - storage allocation 156
- vxfs_inotopath 708
- vxfs_ninode 711
- vxiod I/O kernel threads 48
- vxlsino 708
- vxmake
 - creating cache objects 345
 - creating plexes 611
- vxmend
 - re-enabling plexes 660
- vxmirror
 - configuring VxVM default behavior 612
 - mirroring root disks 685
 - mirroring volumes 612
- vxnotify
 - monitoring configuration changes 604
- vxplex
 - attaching plexes to volumes 611
 - converting plexes to snapshots 378
 - reattaching plexes 660
 - removing mirrors 616
 - removing mirrors of root disk volumes 692
 - removing plexes 616
- vxprint
 - displaying DCO information 385
 - displaying snapshots configured on a cache 374
 - listing spare disks 551
 - listing volumes on boot disks 688
 - verifying if volumes are prepared for instant snapshots 343
 - viewing base minor number 635
- vxrecover
 - recovering plexes 548
 - restarting moved volumes 600
- vxrelayout
 - resuming online relay layout 610
 - reversing direction of online relay layout 611
 - viewing status of online relay layout 610
- vxrelocd
 - hot-relocation daemon 544
 - modifying behavior of 558
 - notifying users other than root 559
 - operation of 545
 - preventing from running 559
 - reducing performance impact of recovery 559
- vxrestore 183
- vxsnap
 - adding snapshot mirrors to volumes 362
 - administering instant snapshots 335
 - backing up multiple volumes 359
 - controlling instant snapshot synchronization 370

vxsnap *(continued)*

- creating a cascaded snapshot hierarchy 363
- creating full-sized instant snapshots 352, 358
- creating linked break-off snapshot volumes 358
- creating space-optimized instant snapshots 350
- displaying information about instant snapshots 368
- dissociating instant snapshots 366
- preparing volumes for instant snapshots 343
- reattaching instant snapshots 364
- reattaching linked third-mirror snapshots 365
- refreshing instant snapshots 364
- removing a snapshot mirror from a volume 363
- restore 335
- restoring volumes 366
- splitting snapshot hierarchies 367

vxspcshow

- discovering device names 98

vxsplittlines

- diagnosing serial split brain condition 653

vxstat

- determining which disks have failed 548

vxtask

- aborting tasks 604
- listing tasks 603
- monitoring online relay layout 610
- monitoring tasks 604
- pausing online relay layout 610
- resuming online relay layout 610
- resuming tasks 604

vxtune

- setting volpagemod_max_memsz 741

vxtune utility 745, 747**vxunreloc**

- listing original disks of hot-relocated subdisks 557
- moving subdisks after hot-relocation 556
- restarting after errors 558
- specifying different offsets for unrelocated subdisks 557
- unrelocating subdisks after hot-relocation 556
- unrelocating subdisks to different disks 556

vxunroot

- removing rootability 692
- unencapsulating the root disk 692

VxVM

- configuration daemon 58
- configuring to create mirrored volumes 612
- dependency on operating system 48
- disk discovery 189

VxVM *(continued)*

- granularity of memory allocation by 734
- maximum number of subdisks per plex 733
- maximum number of volumes 731
- maximum size of memory pool 734
- minimum size of memory pool 736
- objects in 51
- removing disks from 626
- removing disks from control of 663
- rootability 681
- task monitor 601
- types of volume layout 142
- upgrading 622
- upgrading disk group version 622

vxvol

- restarting moved volumes 600
- setting read policy 154
- stopping volumes 663

vxvset

- adding volumes to volume sets 453
- controlling access to raw device nodes 457
- creating volume sets 452
- creating volume sets with raw device access 456
- listing details of volume sets 453
- removing volumes from volume sets 453
- starting volume sets 454
- stopping volume sets 454

W**warning messages**

- Specified region-size is larger than the limit on the system 343

worldwide name identifiers 264**writable Storage Checkpoints** 391**write size** 181**WWN identifiers** 264