

Storage Foundation Administrator's Guide

Windows

7.1

Storage Foundation Administrator's Guide

Document version: 7.1 Rev 0

Last updated: 2016-04-25

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc_feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Overview	23
	About Storage Foundation	23
	Optimized system performance	25
	Software components	25
	Supported software	26
	Advantages of dynamic volumes	26
	Software RAID provided by SFW	27
	What is RAID?	27
	SFW and Virtual Disk Service (VDS)	27
	Choosing RAID levels	28
	Hardware and software RAID	31
Chapter 2	Setup and configuration	32
	Setup and configuration overview	32
	Function overview	33
	Basic and dynamic disks	33
	Basic and dynamic volumes	33
	Disk groups	34
	Dynamic volume components	35
	About the client console for Storage Foundation	36
	Accessing VEA from the Windows Start menu	37
	Accessing VEA through the Windows computer management window	37
	Disconnecting from a server	38
	Understanding the conditions for remote connection	39
	Recommendations for caching-enabled disks	39
	Write caches	39
	HBA caches	40
	Review the Veritas Enterprise Administrator GUI	40
	Configure basic disks (Optional)	42
	About basic disks	42
	Upgrading basic disks to dynamic	42
	About creating dynamic disk groups	43
	Considerations when changing a basic disk to dynamic	43
	Creating a dynamic disk group	44

Making storage available after system startup	47
About creating dynamic volumes	47
Background information	47
Creating a dynamic volume	48
Dynamic volume types	56
Set up a dynamic boot and system volume (Optional)	59
Set desired preferences	61
Hot relocation settings	62
Event notification settings	62
S.M.A.R.T. monitoring for a disk	63
Capacity monitoring for volumes	63
Task throttling	63
Configuration task performance tuning	64
Track alignment settings	64
Protecting your SFW configuration with vxcbt	64
Chapter 3	
Using the GUI to manage your storage	65
VEA Overview	65
Viewing disks and volumes	66
SFW general preferences	66
The URL bar	67
The Toolbar	67
The Perspective bar	68
Systems perspective	69
The left pane	69
The right pane	72
The lower pane	76
The disk view	77
Disk view overview	77
Options for graphical views	78
Expanded view	79
Search	80
Conducting a search	81
Using saved queries	83
Creating search reports	84
Assistant perspective	86
Logs perspective	86
Configuring SFW	87
Control Panel perspective	87
Log settings	88
Rule Manager	88
Configuration backup	88

	Historical statistics	89
	SMTP configuration	89
	System settings	89
	Task settings	94
	Track alignment	95
Chapter 4	Working with disks, partitions, and volumes	97
	Overview	97
	Adding storage	97
	Add a new disk to your system	98
	Add a disk signature to a disk	98
	Add a disk to a dynamic disk group	99
	Create new volumes	101
	Disk tasks	101
	Remove a disk from a dynamic disk group	102
	Remove a disk from the computer	103
	Offline a disk	104
	Update disk information by using rescan	104
	Set disk usage	105
	Evacuate disk	105
	Replace disk	106
	Changing the internal name of a disk	107
	Renaming an enclosure	107
	Work with removable media	108
	Working with disks that support thin provisioning	108
	View disk properties	110
	Veritas Disk ID (VDID)	113
	General Partition/Volume tasks	114
	Delete a volume	115
	Delete a partition or logical drive	115
	Shredding a volume	115
	Refresh drive letter, file system, and partition or volume information	116
	Add, change, or remove a drive letter or path	117
	Renaming a mirror (plex)	118
	Changing the internal name of a volume	119
	Mount a volume at an empty folder (Drive path)	119
	View all drive paths	120
	Format a partition or volume with the file system command	121
	Cancel format	121
	Change file system options on a partition or volume	121
	Set a volume to read only	122

Check partition or volume properties	122
Expand a dynamic volume	124
Expand a partition	128
Safeguarding the expand volume operation in SFW against limitations of NTFS	129
Safeguarding the expand volume operation in SFW against limitations of ReFS	130
Shrink a dynamic volume	130
Dynamic LUN expansion	133
Basic disk and volume tasks	133
What can you do with a basic disk?	134
Converting the partition style of a disk	134
Create primary and extended partitions	135
New logical drives	137
Mark a partition as active	138
Change a basic disk to dynamic	138
Troubleshoot basic disks and volumes	139
Automatic discovery of SSD devices and manual classification as	
SSD	139
Disk media types	139
Supported Solid-State Devices	140
Icon for SSD	140
Enclosure and VDID for automatically discovered On-Host Fusion-IO disks	140
Enclosure and VDID for automatically discovered On-Host Intel disks	141
Enclosure and VDID for automatically discovered Violin disks	141
Classifying disks as SSD	142
Limitations for classifying SSD devices	144
Volume Manager space allocation is SSD aware	144
Setting an SSD plex as a preferred plex to improve I/O performance	145
Chapter 5 Dealing with disk groups	146
Disk groups overview	146
Basic and dynamic disk groups	147
Moving dynamic disk groups between computers	147
Primary and secondary dynamic disk groups	147
Cluster disks and cluster dynamic disk groups	148
Private dynamic disk group protection	148
Delete a dynamic disk group	149

	Recommended method for deleting a disk group	149
	Alternative method for deleting a disk group	150
	Upgrading the dynamic disk group version	150
	Converting a Microsoft Disk Management Disk Group	152
	Importing a dynamic disk group to a cluster disk group	153
	Rename a dynamic disk group	154
	Detaching and attaching dynamic disks	154
	Detaching dynamic disks	155
	Attaching dynamic disks	156
	Importing and deporting dynamic disk groups	157
	Deport a dynamic disk group	157
	Import a dynamic disk group	159
	Partitioned shared storage with private dynamic disk group protection	162
	Create a new dynamic disk group with private dynamic disk group protection	163
	Add private dynamic disk group protection to an existing dynamic disk group	164
	Remove private dynamic disk group protection from a dynamic disk group	165
	Dynamic disk group properties	165
	Troubleshooting problems with dynamic disk groups	167
Chapter 6	Fast failover in clustered environments	168
	What is fast failover	168
	SFW changes for fast failover	169
	Enabling fast failover for disk groups in a VCS environment	171
	Enabling fast failover for disk groups in a Microsoft Failover Clustering environment	172
	Limitations	175
	Configuration and functionality not supported	175
Chapter 7	iSCSI SAN support	176
	Overview of iSCSI SAN support	176
	About iSNS servers	177
	Configuring an iSCSI SAN with SFW	177
	Considering prerequisites	178
	Setting up an iSCSI SAN	178
	Using SFW vxdg lateststart for iSCSI	179
	Managing an iSCSI SAN with SFW	180
	Managing with the VEA Console	180
	About the iSCSI initiators node	180

	About multiple sessions	183
	About the iSCSI arrays node	183
	About the iSNS server node	185
Chapter 8	Settings for monitoring objects	188
	Overview	188
	Event monitoring and notification	188
	Alert log	189
	Severity levels	189
	Log configuration	189
	Event notification	190
	Accessing the Rule Manager utility	190
	Preliminary setup	190
	Steps for creating rules	191
	Accessing the SFW SNMP MIB files	193
	Disk monitoring	193
	Capacity monitoring	195
	Configuring Automatic volume growth	197
	Enabling Automatic Volume Growth in a non-clustered and clustered environment	199
	Configuring the Veritas Scheduler Service on a clustered environment	200
	Automatic volume growth script for volume capacity monitoring	201
	SMTP configuration for email notification	204
Chapter 9	Standard features for adding fault tolerance	206
	Fault tolerance overview	206
	Mirroring	206
	About mirroring	207
	Two methods for creating a mirror	208
	Break or remove a mirror	210
	Break mirror	210
	Remove mirror	211
	Set the mirrored volume read policy	212
	Mirror a dynamic boot or system volume	214
	Reactivate a mirrored volume	214
	Repair a mirrored volume	215
	RAID-5 and RAID 0+1	215
	RAID-5	215
	RAID 0+1	216
	Hot relocation	216

	Hot relocation overview	217
	Hot relocation setup summary	218
	Hot relocation mode	218
	Customize target disks with set disk usage	219
	Undo hot relocation	220
	Clear hot relocation information	221
	Dirty region logging (DRL) and RAID-5 logging	222
	Dirty region logging for mirrored volumes	222
	RAID-5 logging	223
	Add a log	223
	Remove a log	224
	Dynamic relayout	225
	Examples of dynamic relayout	226
	Detailed steps for dynamic relayout	227
Chapter 10	Performance tuning	229
	Overview	229
	Statistics overview	229
	Real-time statistics collection and display	231
	Online monitoring window features	231
	Setup steps for real-time statistics	233
	Historical statistics collection and graphing	237
	Features of the graphing window	237
	Setup steps for historical statistics	238
	Subdisk move, split, and join	241
	Purpose of the subdisk commands	242
	Subdisk context menu	243
	Subdisk properties	243
	Moving subdisks	244
	Splitting subdisks	245
	Joining subdisks	246
	Command line commands for statistics	247
	SmartMove	247
Chapter 11	FlashSnap	248
	FlashSnap overview	248
	FlashSnap components	250
	Summary of the FlashSnap procedure	251
	FastResync	253
	FastResync limitations	254
	Resynchronization	254
	Reconnection	255

Enabling and disabling FastResync	255
Disk Change Object (DCO) volume	256
Snapshot commands	259
Using the snapshot procedure	260
Prepare	261
Snap shot	263
Snap back	265
Snap clear	266
Snap abort	267
Using the Snapshot Scheduler wizard	268
Dynamic Disk Group Split and Join	272
About Dynamic Disk Group Split and Join	273
Dynamic disk group split	274
Recovery for the split command	277
Dynamic disk group join	279
Using Dynamic Disk Group Split and Join with a cluster on shared storage	282
Limitations when using dynamic disk group split and join with Volume Replicator	285
Dynamic Disk Group Split and Join troubleshooting tips	285
CLI FlashSnap commands	286
Fast File Resync	287
Overview	287
Volume Shadow Copy Service (VSS)	289
Overview	289
VSS components	290
Overview of VSS process for SFW snapshots	292
VSS snapshot wizard for Microsoft Exchange and Microsoft SQL	293
Using the VSS snapshot wizards with Microsoft Exchange	295
Using the VSS Snapshot wizard	296
Using the VSS Snapback wizard	300
Using the VSS Snapshot Scheduler wizard	301
Using the VSS Restore Wizard	309
Refreshing the snapshot set	319
Supporting replication in Exchange 2007	320
Supporting replication in Exchange 2010	323
Using the VSS snapshot wizards with Enterprise Vault	325
Using the Enterprise Vault Snapshot wizard	327
Using the Enterprise Vault Snapback wizard	333
Using the VSS Enterprise Vault Snapshot Scheduler wizard	334
Using the Enterprise Vault Restore Wizard	339
Using the VSS snapshot wizards with Microsoft SQL	342

Using the VSS Snapshot wizard	343
Using the VSS Snapback wizard	347
Using the VSS Snapshot Scheduler wizard	347
Using the VSS Restore Wizard	355
Copy on Write (COW)	359
COW overview	359
SFW VSS COW snapshot process	360
Managing Shadow Storage	360
Using the VSS COW Snapshot wizard	361
Using the VSS COW Snapshot Scheduler wizard	362
Using the VSS COW Restore Wizard	366
Additional information about COW snapshots	367
Using the VSS COW snapshot wizards with Microsoft Exchange	368
Using the VSS COW Snapshot wizard	368
Using the VSS COW Snapshot Scheduler wizard	370
Using the VSS COW snapshot wizards with Microsoft SQL	374
Using the VSS COW Snapshot wizard	375
Using the VSS COW Snapshot Scheduler wizard	376
Chapter 12	
Configuring data caching with SmartIO	382
About SmartIO	382
Typical deployment scenarios	383
Cluster deployment in a Hyper-V environment	383
Cluster deployment in a VMware environment- No vMotion support	384
Cluster deployment in a VMware environment- with vMotion support	385
How SmartIO works	386
SmartIO benefits	387
SmartIO limitations	388
About cache area	388
Attributes of the cache area	388
About SmartIO caching support	389
Configuring SmartIO	390
Creating and administering cache area using VEA	391
Creating and administering cache area through CLI	394
Frequently asked questions about SmartIO	395
How to configure a volume to use a non-default cache area?	395
What is write-through I/O caching?	396
Does SmartIO with SFW support write caching?	396
Are there any logs that I can refer to, if caching fails for a particular volume?	396

- I have deleted a cache area, but the disk is still present in the Cachepool. How can I remove it from the Cachepool? 396
- Is the VxVM cached data persistent? 396
- Is an application's performance affected if the cache device becomes inaccessible while caching is enabled? 397
- Are there any tools available to measure SmartIO performance? 397
- Will there be a performance drop after vMotion? 397
- Will the cache area recreation fail, if the SmartDisk assigned has insufficient space? 397
- A cache area recreation is in process in a VMware environment with vMotion, does it affect the sfcache operations? 397
- Does SmartIO continue to use the previous cache area if the VM is moved back to the previous host? 397
- How does SmartIO behave if the vxsvc service fails during vMotion? 398

Chapter 13 Dynamic Multi-Pathing 399

- DMP overview 399
- Major features of Dynamic Multi-Pathing 400
- Active/Active and Active/Passive settings 401
 - Active/Active and Active/Passive settings in a cluster environment 404
- DMP DSMs 407
 - DMP DSMs menus 408
 - Add and remove paths 412
 - Specify load balancing settings and the primary path 413
 - Specifying control timer settings for an array 420
 - View array, disk, and path status 421

Chapter 14 Configuring Cluster Volume Manager (CVM) 424

- Overview 424
- Configuring a CVM cluster 425
 - Preparing SFW for Microsoft Failover Clustering 426
 - Creating cluster-shared disk groups and volumes 429
 - Configuring a role for volume-level failover 432
- Administering CVM 433
 - Configuring CVM links for multi-subnet cluster networks 433
 - Access modes for cluster-shared volumes 434
 - Storage disconnectivity and CVM disk detach policy 437
- Unconfiguring a CVM cluster 443
 - Unconfiguring CVM using ffconfigcli 444

	Command shipping	445
	Operations supported by command shipping	446
	Operations supported for running where volume is online	446
	About I/O Fencing	447
	About preventing data corruption with I/O Fencing	447
	About I/O Fencing operations	448
	About I/O Fencing components	448
Chapter 15	Administering site-aware allocation for campus clusters	449
	Overview	449
	Pre-requisites for enabling site-aware allocation on a campus cluster setup	451
	About the site-aware read policy	452
	Administering hosts for site-based allocation	452
	Administering hosts for site-aware allocation through GUI	453
	Administering hosts for site-aware allocation through CLI	453
	Administering disks for site-based allocation	454
	Administering disks for sites through GUI	454
	Administering disks for sites through CLI	456
	Administering volumes for site-based allocation	456
	Setting or resetting volume Sitetype property for site allocation	457
	Creating a dynamic volume	459
	Administering mirrored plexes for site-based allocation	460
	Adding a mirror to a site-based volume through GUI	460
	Adding a mirror to a site-based volume through CLI	465
Chapter 16	SFW for Hyper-V virtual machines	466
	Introduction to Storage Foundation solutions for Hyper-V environments	466
	Advantages of running Storage Foundation in the Hyper-V parent	466
	How Storage Foundation manages storage for virtual machines	467
	Related documentation on Storage Foundation	468
	Live migration support for SFW dynamic disk group	469
	About implementing Hyper-V virtual machine live migration on SFW storage	470
	Tasks for deploying live migration support for Hyper-V virtual machines	470
	Installing Windows Server	472

Preparing the host machines	472
Installing the SFW option for Microsoft failover cluster option	473
Using the SFW Configuration Wizard for Microsoft Failover Cluster for Hyper-V live migration support	473
Configuring the SFW storage	474
Creating a virtual machine service group	475
Setting the dependency of the virtual machine on the VMDg resource	476
Administering storage migration for SFW and Hyper-V virtual machine volumes	477
About storage migration	478
About performance tunables for storage migration	478
Setting performance tunables for storage migration	478
About performing online storage migration	480
Storage migration limitations	481
About changing the layout while performing volume migration	481
Migrating volumes belonging to SFW dynamic disk groups	481
Migrating volumes belonging to Hyper-V virtual machines	483
Migrating data from SFW dynamic disks of one enclosure to another	486
Converting your existing Hyper-V configuration to live migration supported configuration	488
Optional Storage Foundation features for Hyper-V environments	490
About using optional Storage Foundation features in the Hyper-V parent	491
Dynamic Multi-Pathing for the virtual environment	491
Replicating virtual machines	492
Virtual machine volume snapshots	493
Campus clusters	494
Chapter 17 Microsoft Failover Clustering support	495
Overview	495
Enabling Microsoft Failover Clustering after InfoScale Storage is already installed	496
Configuring Microsoft Failover Clustering	498
Create a cluster dynamic disk group	500
Additional considerations for SFW Microsoft Failover Clustering support	503

Chapter 18	Configuring a quorum in a Microsoft Failover Cluster	507
	Overview	507
	SFW support for dynamic quorum resource	508
	Configuring Failover Cluster quorum using SFW	508
	Configuring cluster quorum settings	510
	Microsoft Failover Cluster quorum arbitration time settings	511
Chapter 19	Implementing disaster recovery with Volume Replicator	513
	Overview	513
	Volume Replicator feature highlights	514
	How Volume Replicator works	514
	Summary of the steps for setting up a disaster recovery solution with Volume Replicator	516
	Volume Replicator terms	518
	Replicated Volume Group (RVG)	518
	Replicated Data Set (RDS)	518
	Replicator Log	518
	Replication Modes	519
Chapter 20	Troubleshooting and recovery	520
	Overview	520
	Using disk and volume status information	521
	Disk status descriptions	521
	Volume status descriptions	524
	SFW error symbols	527
	Resolving common problem situations	528
	Bring an offline dynamic disk back to an imported state	529
	Bring a basic disk back to an online state	530
	Remove a disk from the computer	531
	Bring a foreign disk back to an online state	531
	Bring a basic volume back to a healthy state	532
	Bring a dynamic volume back to a healthy state	533
	Repair a volume with degraded data after moving disks between computers	534
	Deal with a provider error on startup	535
	Commands or procedures used in troubleshooting and recovery	536
	Refresh command	536
	Rescan command	536
	Replace disk command	538

Merge foreign disk command	538
Reactivate disk command	540
Reactivate volume command	541
Repair volume command for dynamic RAID-5 volumes	542
Repair volume command for dynamic mirrored volumes	542
Starting and stopping the Storage Foundation Service	543
Accessing the CLI history	544
Additional troubleshooting issues	544
Disk issues	545
Volume issues	546
Disk group issues	549
Connection issues	552
Issues related to boot or restart	553
Cluster issues	555
Dynamic Multi-Pathing issues	558
vxsnap issues	558
Other issues	560
CVM issues	562
After a storage disconnect, unable to bring volume resources online on the CVM cluster nodes	562
Error may occur while unconfiguring a node from CVM cluster	563
Shutdown of all the nodes except one causes CVM to hang	563
Sometimes, CSDG Deport causes Master node to hang due to IRP getting stuck in QLogic driver	563
Unknown disk groups seen on nodes after splitting a cluster-shared disk group into cluster disk groups from Slave node	564
In some cases, missing disks are seen on target Secondary dynamic disk groups after splitting a cluster-shared disk group from Slave node	564
Cannot stop vxsvc if SFW resources are online on the node	564
Cluster-shared volume fails to come online on Slave if a stale CSDG of the same name is present on it	564
CVM does not start if all cluster nodes are shut down and then any of the nodes are not restarted	565
Incorrect errors shown while creating a CSDG if Volume Manager Shared Volume is not registered	565
After splitting or joining disk group having mirrored volume with DRL, VEA GUI shows incorrect volume file system if volumes move to another disk group	565
Enclosure-level storage migration fails, but adds disks if a cluster-shared volume is offline	566

Volume Manager Shared Volume resource fails to come online
or cannot be deleted from Failover Cluster Manager 566

Sometimes, source cluster-shared volumes are missing after
joining two cluster-shared disk groups 567

If private CVM links are removed, then nodes may remain out of
cluster after network reconnect 567

Format dialog box appears after storage disconnect 567

Volume Manager Shared Volume resources fail to come online
on failover nodes if VxSVC is stopped before stopping
clussvc 568

One or more nodes have invalid configuration or are not running
or reachable 568

After node crash or network disconnect, volume resources failover
to other node but the drive letters are left behind mounted on
the failing node even after it joins cluster successfully 569

Shutdown of Master node in a CVM cluster makes the Slave nodes
to hang in “Joining” state while joining to new Master 570

CVM stops if Microsoft Failover Clustering and CVM cluster
networks are not in sync because of multiple, independent
network failures or disconnect 570

Restarting CVM 570

Administering CVM using the CLI 571

Tuning the VDS software provider logging 571

Appendix A Command line interface 573

Overview of the command line interface 574

SFW commands available from the command line 574

Command line log 576

Conventions for command line syntax 577

vxclustadm 580

vxclustadm startnode 580

vxclustadm stopnode 580

vxclustadm nodestate 580

vxclustadm nidmap 581

vxvol 581

vxvol volinfo 582

vxvol repair 582

vxvol set fastresync=on|off 583

vxvol reactivate 584

vxvol rdpol prefer 585

vxvol rdpol round 586

vxvol growfs 587

vxvol access	588
vxvol reclaim	589
vxvol reclaimcancel	590
vxvol capacitymonitor	590
vxvol autogrow	591
vxvg	593
vxvg list	595
vxvg list (CSDG)	596
vxvg dginfo	596
vxvg dginfo (CSDG)	597
vxvg init	599
vxvg init (CSDG)	601
vxvg adddisk	602
vxvg rmdisk	603
vxvg import	604
vxvg import (CSDG)	606
vxvg deport	608
vxvg destroy	608
vxvg protect	608
vxvg release	609
vxvg upgrade	610
vxvg repldisk	611
vxvg split	612
vxvg recover	615
vxvg join	616
vxvg reclaim	618
vxvg reclaimcancel	618
vxvg latestart	619
vxvg refreshff	620
vxclus	621
vxclus enable	622
vxclus disable	623
vxclus cleanup	623
vxclus UseSystemBus ON	623
vxclus UseSystemBus OFF	624
vxdisk	624
vxdisk diskinfo	625
vxdisk attach	626
vxdisk detach	627
vxdisk list	628
vxdisk list (CSDG)	628
vxdisk merge	629
vxdisk reactivate	630

vxdisk set hotreloc	630
vxdisk hotrelocinfo	632
vxdisk sig	632
vxdisk set track	633
vxdisk trackaligninfo	634
vxdisk ping	634
vxdisk reclaim	635
vxdisk reclaimcancel	636
vxdisk setsite	637
vxdisk rmsite	637
vxassist	638
vxassist make	639
vxassist growby	644
vxassist querymax	647
vxassist shrinkby	648
vxassist shrinkabort	649
vxassist mirror	650
vxassist break	655
vxassist remove	656
vxassist delete	659
vxassist shred	659
vxassist addlog	661
vxassist online (read/write)	663
vxassist offline	663
vxassist prepare	664
vxassist snapshot	668
vxassist snapback	669
vxassist snapclear	670
vxassist snapabort	672
vxassist rescan	673
vxassist refresh	673
vxassist resetbus	674
vxassist version	674
vxassist (Windows-specific)	674
vxassist assign	674
vxassist unassign	675
vxassist create_part	676
vxassist create_ext	676
vxassist delete_ext	676
vxassist create_drive	677
vxassist taghost	677
vxassist showtag	677
vxassist untaghost	678

vxassist set sitetype	678
vxassist resetsitetype	678
vxevac	678
vxsd	679
vxsd mv	680
vxsd split	680
vxsd join	681
vxstat	681
vxtask	684
vxedit	687
vxedit set comment	687
vxedit set alias	688
vxedit rename	688
vxunreloc	689
vxdmpadm	690
vxdmpadm dsminfo	691
vxdmpadm arrayinfo	691
vxdmpadm deviceinfo	692
vxdmpadm pathinfo	692
vxdmpadm arrayperf	693
vxdmpadm deviceperf	693
vxdmpadm pathperf	694
vxdmpadm allperf	694
vxdmpadm iostat	695
vxdmpadm cleardeviceperf	696
vxdmpadm cleararrayperf	697
vxdmpadm clearallperf	697
vxdmpadm setdmscsi3	697
vxdmpadm setarrayscsi3	698
vxdmpadm setattr dsm	698
vxdmpadm setattr array	700
vxdmpadm setattr device	702
vxdmpadm setattr path	703
vxdmpadm set isislog	705
vxdmpadm rescan	705
vxdmpadm disk list	705
vxdmpadm getdsmattrib	706
vxdmpadm getmpioparam	706
vxdmpadm setmpioparam	707
vxabr	708
vxabr backup	708
vxabr restore	709
vxabr write_signature	710

vxsnap	710
vxsnap prepare	712
vxsnap create	716
vxsnap reattach	727
vxsnap restore	731
vxsnap refresh	740
vxsnap diffarea	741
vxsnap delete	742
vxfsync	743
vxscrub	743
vxscrub [-a -b -l -v]	744
vxscrub -p	745
vxscrub -forcepurge	745
vxverify	746
vxprint	747
sfcache	749
sfcache create	751
sfcache delete	751
sfcache enable	751
sfcache disable	752
sfcache online	752
sfcache offline	752
sfcache growby	753
sfcache shrinkby	753
sfcache set	753
sfcache list	753
sfcache stat	754
Tuning SFW	754
Displaying the tunable values	755
Setting the tunable values	755
Appendix B VDID details for arrays	757
Format of VDID strings for disks belonging to various arrays	757
Glossary	760

Overview

This chapter includes the following topics:

- [About Storage Foundation](#)
- [Software components](#)
- [Supported software](#)
- [Advantages of dynamic volumes](#)
- [Software RAID provided by SFW](#)

About Storage Foundation

Storage Foundation is available as part of InfoScale Storage and InfoScale Enterprise.

SFW and SFW HA have the same functions except for the feature of clustering support. SFW HA adds Cluster Server (VCS) support. SFW retains the cluster option for Failover Cluster feature.

InfoScale Storage and InfoScale Enterprise differ in the clustering support available. InfoScale Enterprise provides clustering support using VCS. Whereas, InfoScale Storage provides clustering support using Microsoft Failover Cluster.

Storage Foundation provides a comprehensive solution to storage management in an integrated graphical view. SFW lets you configure and manage the local and the remote storage that is attached to your system while your system remains online.

With Storage Foundation, you can do the following:

- Manage domain-wide storage from a single console.
- Create and configure software RAID while systems remain online and available.

- Optimize storage system performance by online capacity monitoring and storage load balancing.
- Provides the ability to configure and manage different volume layouts: Concatenated, striped, mirrored, mirrored striped, and RAID-5 volumes. Supports up to 32-way mirrors on a mirrored volume.
- Provides automatic detection of failed disks and the ability to repair fault-tolerant volumes on those disks without interrupting applications using the volumes.
- Provides a storage migration path with the ability to upgrade existing partitions to the new "dynamic" volumes (volumes based on Storage Foundation technology).
- Supports online extending of all volume types.
- Supports a cluster environment that runs under Microsoft Cluster Service software.
- Supports mounting a volume without requiring the use of a drive letter.
- Supports moving of storage between computers with the Import and Deport Dynamic Disk Group functions.
- Provides a command-line support for administrators who prefer this method over a graphical user interface (GUI).
 For a Server Core installation of Windows Server, only the SFW command line interface (CLI) is available.
- Provides a command-line support for administrators to manually pause or resume the tasks running.
 You can pause or resume a task only on a node where the task is running. In case of a CVM cluster, if a command is executed from a slave and a task begins on a master node, then the task can be paused only on the master node. The details of the tasks that are paused or resumed are shown in Alerts. The following tasks can be paused and resumed:
 - Online storage migration
 - Volume shred
 - Mirroring
 - Subdisk move

Note: While resuming, the tasks resume from the current state of completion. However, if the system reboots or the resource fails over when the task is paused, the task is reinitialized.

Optimized system performance

Performance degradation occurs when a program or device uses too much disk I/O. By monitoring a system's I/O, you can find the problem areas and eliminate areas of high activity ("hot spots"), thus ensuring smooth operation. SFW's Performance Monitoring utility lets you continuously monitor and tune system performance, monitor I/O statistics, perform hot spot detection, and make adjustments.

Storage Foundation can help improve overall disk and system performance in several ways, some of which are as follows:

- I/O statistics are used to identify high-traffic areas, known as "hot spots." You can use the Move Subdisk command to resolve these hot spots online.
- Data is assigned to physical drives to evenly balance the I/O load among the disk drives. This is known as storage load balancing.
- Event logging of errors and important information is provided.
- Tasks can be paused and resumed to provide additional bandwidth and CPU resources for the priority tasks.

Software components

Storage Foundation includes the following components:

Client software	The client software includes the console and runs on the supported Windows server and client operating systems. The client enables you to configure and manage the storage that is attached to local and to remote hosts.
Server software	The server software, which runs on a managed server, is the common repository for all storage objects.
Providers	<p>The providers run on a managed server. Providers are similar to drivers. Each provider manages a specific hardware or software storage component.</p> <p>The providers discover the existing physical and logical entities and store that information in the Storage Foundation distributed database. Providers update the database whenever there is a change in the physical or the logical entities present in the hardware or software.</p>

Supported software

For the latest information on the supported hardware and software for the InfoScale products, see the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:

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

Advantages of dynamic volumes

One of the main advantages of dynamic volumes is that they can span across multiple disks. You can get better use of your available disk space by combining areas of unallocated space in the volumes that go across multiple disks. Multi-disk volumes also improve disk performance by letting more than one disk drive read and write data. Finally, multi-disk volumes let you use RAID technology to make volumes fault tolerant.

When you upgrade a disk so that it can contain dynamic volumes, Storage Foundation writes a small database onto the upgraded disk that contains volume management information about the disk and all other dynamic disks in its disk group. This database gives a dynamic disk additional advantages, which are as follows:

- Users can create dynamic volumes without restarting the system.
- Volumes can be accessed continuously, even across disk reconfiguration events.
- All dynamic volumes with an NTFS or ReFS file system can be extended online. You can select where to extend the volume, or Storage Foundation assigns space optimally on the disks available within the volume's dynamic disk group.
- Disk configuration information is contained on the disk, not in the registry. The disk configuration information is also replicated across all dynamic disks in the disk group so that the failure of one disk does not obstruct access to data on other disks.
- Dynamic disks provide volume migration, which is the ability to move a disk or disks containing a volume or volumes from one system to another system without loss of data.
- Dynamic disks let you move portions of volumes (subdisks) between disks on a single computer system to optimize performance.

Software RAID provided by SFW

To make the best use of the dynamic volumes that SFW provides, you need an understanding of the RAID levels available with the program and what uses are appropriate for each RAID level.

See [“What is RAID?”](#) on page 27.

See [“SFW and Virtual Disk Service \(VDS\)”](#) on page 27.

See [“Choosing RAID levels”](#) on page 28.

See [“Hardware and software RAID”](#) on page 31.

What is RAID?

RAID (Redundant Array of Independent Disks) is a collection of specifications that describe a system for ensuring the reliability and stability of data that is stored on large disk subsystems.

RAID provides the following general benefits:

- Increased storage system reliability
- Improved I/O performance
- Fault tolerance

SFW and Virtual Disk Service (VDS)

Virtual Disk Service (VDS) provides a standard way to manage multi-vendor storage devices through a common storage interface. VDS focuses on the management of logical volumes through both software and hardware RAID. VDS supports the Disk Management user interface and two scriptable command line interfaces, DISKPART and DISKRAID.

DISKPART implements software RAID procedures. It deals with the creation, extension, and deletion of logical volumes or partitions on both basic disks and dynamic disks. DISKPART is the command-line version of Disk Management.

DISKRAID implements hardware RAID procedures. It is used with vendor-specific hardware arrays to configure virtual disks (logical disks that span two or more independent physical disks).

In order for the common storage interface to work, participating hardware and software vendors must prepare their providers for the VDS software. Each hardware or software provider translates VDS standard APIs into instructions specific to the related storage device or software RAID procedures.

Storage administrators can use the scriptable command line utilities to create scripts to automate storage procedures across multi-vendor devices. The administrator creates the scripts using standard commands from the CLI, and therefore does not have to spend extra time to learn the specific management utilities that come with each hardware device or software RAID program. Once the script is set up, it can automate time-consuming procedures that had been done manually, such as adding storage to the network.

For Windows Server operating systems, the native disk and volume manager is Microsoft Disk Management. On a Windows Server OS, SFW does not take the place of Microsoft Disk Management. SFW provides all the functionality as if it were installed on a Windows Server OS with the added benefit that Microsoft Disk Management is still available. Uninstalling SFW components does not affect Microsoft Disk Management.

On a Windows Server operating system, SFW can be accessed with the Microsoft Disk Management GUI.

Note: On Windows Server operating systems, Microsoft Disk Management does not support objects that SFW creates, such as dynamic disks or dynamic volumes.

At this time, the advanced functionality of Storage Foundation is not available through the Microsoft CLI VDS interfaces.


Choosing RAID levels

The following illustrations show how RAID levels (RAID-0, RAID-1, RAID-5, and RAID 0+1) provide various levels of redundancy and performance, compared with a simple volume on a single disk. From these, you can decide what choices are applicable to your particular situation.

Concatenated

The following figure illustrates a concatenated volume layout.

Figure 1-1 Concatenated volume layout




Concatenated volume layout offers the following:

- Concatenates n disks as one large virtual disk with a capacity of n disks.
- Data fills up the first disk before it is written to the second disk.
- Not fault tolerant. When a disk fails, the large virtual disk fails.
- No performance gain.

RAID level 0 (Striping)

The following figure illustrates a RAID level 0 volume layout.

Figure 1-2 RAID level 0 volume layout




RAID level 0 volume layout offers the following:

- Groups n disks as one large virtual disk with a capacity of n disks.
- Data is stored to the disks alternately.
- Not fault tolerant. When a disk fails, the large virtual disk fails.
- Better read and write performance than both simple volumes and all other RAID levels.

RAID level 1 (Mirroring)

The following figure illustrates a RAID level 1 volume layout (mirroring).

Figure 1-3 Mirrored volume




RAID level 1 (mirroring) volume layout offers the following:

- Groups two or more disks as one virtual disk with the capacity of a single disk.
- Data is replicated on each disk, providing data redundancy.
- When a disk fails, the virtual disk still works. The data is read from the surviving disk(s).
- Better read performance, but slightly slower write performance.
- Because data is duplicated on two or more disks, RAID-1 is more "expensive" in terms of disk space.

RAID level 5 (Striping with distributed parity)

The following figure illustrates a RAID level 5 volume layout (striping with distributed parity).

Figure 1-4 RAID-5 volume layout



RAID-5 (striping with distributed parity) volume layout offers the following:


- Groups n the disks as one large virtual disk with a capacity of $(n-1)$ disks.
- Redundant information (parity) is alternately stored on all disks.

- When a disk fails, the virtual disk still works. The data is reconstructed from the surviving disks and the parity alternately stored on the disks. You need to replace the failed disk to restore fault tolerance.
- Better read performance, but slower write performance.

RAID level 0+1 (Mirrored striped volumes)

The following figure illustrates a RAID level 0+1 volume layout (mirrored striped volumes).

Figure 1-5 RAID level 0+1 volume layout



RAID level 0+1 (mirrored striped volumes) volume layout offers the following:

- Stripes the data across two disks and mirrors it.
- When a disk fails, the data is read from the surviving mirror.
- Better read performance and write performance than RAID level 1 alone.
- Provides both redundancy (with mirroring) and performance (with striping).

Hardware and software RAID

RAID functions can be implemented with either hardware RAID (a specific RAID controller) or software RAID (software without the specific hardware).

Software RAID can be configured on top of the hardware RAID volumes and provide a powerful variety of features.

Setup and configuration

This chapter includes the following topics:

- [Setup and configuration overview](#)
- [Function overview](#)
- [About the client console for Storage Foundation](#)
- [Recommendations for caching-enabled disks](#)
- [Review the Veritas Enterprise Administrator GUI](#)
- [Configure basic disks \(Optional\)](#)
- [About creating dynamic disk groups](#)
- [About creating dynamic volumes](#)
- [Set up a dynamic boot and system volume \(Optional\)](#)
- [Set desired preferences](#)
- [Protecting your SFW configuration with vxabr](#)

Setup and configuration overview

This section provides details about the tasks that are required to configure storage, using Storage Foundation. It also describes how to work with disk groups that are created with an earlier versions of the software, Volume Replicator.

Storage Foundation lets you optimize your storage through the use of dynamic disk groups and dynamic volumes.

The main configuration tasks that you need to do are as follows:

- Create one or more disk groups.

- Create dynamic volumes for each disk group.
- Set up a dynamic boot and / or system volume (optional).
- Configure options.
- Set preferences.

Function overview

Before you can configure Storage Foundation, you need an understanding of the main concepts in the program.

See [“Basic and dynamic disks”](#) on page 33.

See [“Basic and dynamic volumes”](#) on page 33.

See [“Disk groups”](#) on page 34.

See [“Dynamic volume components”](#) on page 35.

Basic and dynamic disks

Physical disks under SFW's control are either basic or dynamic. Disks that are configured on a fresh system are defined as basic disks. You can upgrade them to dynamic disks by making the disks part of a dynamic disk group.

Basic and Dynamic disk details are as follows:

Basic disks	Basic disks adhere to the partition-oriented scheme of Windows NT, Windows 95/98, and MS-DOS. Basic disks can also contain RAID volumes that were created in NT Disk Administrator. In addition, CD-ROMs and other removable media are considered basic disks.
Dynamic Disks	Dynamic disks can contain dynamic volumes that are created with Storage Foundation. The five dynamic volume types are concatenated (includes simple and spanned volumes), mirrored, striped, RAID-5, and mirrored striped (RAID 0+1). On a dynamic disk, space is organized through dynamic volumes rather than partitions.

Because a dynamic disk does not have the partitioning scheme used by Windows NT, Windows 95/98, and MS-DOS, you cannot access dynamic disks through those operating systems.

Basic and dynamic volumes

A volume is a logical entity that is made up of a portion or portions of one or more physical disks. You can format a volume with a file system and can access it with

a drive letter or a mount path. Like disks, volumes can be basic or dynamic. Volumes are defined as basic or dynamic, depending on whether they reside on a basic or dynamic disk.

Basic and Dynamic volume details are as follows:

Basic volumes	Basic volumes can be primary or extended partitions, simple logical drives that reside on extended partitions, or RAID volumes that were originally created in Windows NT Disk Administrator.
Dynamic volumes	<p>Dynamic volumes are volumes created on dynamic disks by using Storage Foundation. You can create any number of dynamic volumes in unallocated space on one or more disks.</p> <p>The volumes are created online without the need for restarting (except a boot or system volume). Each volume can have a FAT, FAT32, NTFS, or ReFS file system.</p> <p>SFW's volumes are internally organized according to established RAID levels.</p> <p>See "Software RAID provided by SFW" on page 27.</p> <p>You can do online extending of all dynamic volume types, and you can add up to 32 mirrors to any volume type, except RAID-5.</p> <p>Windows Server accesses the dynamic volumes that SFW creates in the same way that it accesses the physical partitions. The volumes are identified by a drive letter or mount point.</p>

Although you cannot create new NT Disk Administrator-type volumes on a basic disk, you can upgrade them to dynamic. Then these volumes have the full capabilities of other SFW dynamic volumes.

About dynamic volumes in SFW Basic

As in Storage Foundation, dynamic volumes can be created in Storage Foundation Basic. However, SFW Basic is limited to a maximum of four dynamic volumes. Upgrading SFW Basic to SFW lets you create any number of dynamic volumes.

Disk groups

All basic and dynamic volumes should be in a disk group. There is only one disk group for basic disks, known as the Basic disk group (or BasicGroup in the VEA GUI). There can be one or more dynamic disk groups. A dynamic disk group contains dynamic disks managed by SFW.

When you move disks between computers, you use the Import and Deport Dynamic Disk Group commands. The disk group for clusters is called a cluster disk group.

See “[Disk groups overview](#)” on page 146.

See “[Create a cluster dynamic disk group](#)” on page 500.

On Windows Server operating systems, there are two different types of dynamic disk groups, the Microsoft Disk Management Disk Group and the SFW dynamic disk group. The Microsoft Disk Management Disk Group is a disk group that can be created by Microsoft Disk Management. In the VEA GUI, this type of disk group is distinguished with an icon containing the character "M". In addition to the SFW disk groups, SFW can also create and manage the Microsoft Disk Management Disk Group.

Dynamic volume components

The components or virtual storage objects that are involved in a dynamic volume include its disk groups, the dynamic disks it is associated with, and its plexes, columns, and subdisks. The terms “dynamic disk” and “dynamic disk group” have been defined in the preceding sections.

The definitions for plexes, columns, and subdisks are given below:

Plexes	A plex refers to an instance of the volume. Mirrored volumes have two or more plexes. All other volumes have one plex.
Columns	A column refers to an area on the disk where all or a portion of the volume resides. Striped, RAID-5, and mirrored striped (RAID 0+1) volumes contain multiple columns. Other volume types contain one column.
Subdisks	<p>Subdisks identify the disk location and size of each contiguous disk area in the volume. Subdisks are the basic units in which Storage Foundation allocates disk space in a volume. An SFW disk can be divided into one or more subdisks. Each subdisk represents a specific portion of a volume on a dynamic disk.</p> <p>An SFW disk may contain multiple subdisks, but subdisks cannot overlap or share the same portions of disk space. Any SFW disk space that is not part of a subdisk is considered to be unallocated space, which can be used to create new volumes or to extend existing volumes.</p> <p>You can move subdisks to improve disk performance.</p>

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support operations on subdisks.

About the client console for Storage Foundation

A client console (user interface), Veritas Enterprise Administrator (VEA) is available to perform the operations that are required to configure and manage storage. With the product installation, VEA is installed by default.

VEA GUI can be accessed through either the Start menu or the Computer Management window of the server or the client computer you use. Once you access the VEA GUI, you then need to connect to one or more servers that you want SFW to manage.

Note: You cannot install the older version of VEA (client-only) to manage storage using the latest version of the product.

This section describes how to use these two methods to access VEA and to connect to a server.

Note: To access and perform operations within SFW, you must have administrator rights on the computer you want to connect to. Not having administrator rights only lets you view information in SFW. If User Access Control (UAC) is enabled, run the program or commands in the "Run as administrator" mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

Note: If User Access Control (UAC) is enabled on Windows Server operating systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set `Write` permission for the guest user using the Security tab of the `Veritas` folder's properties.

See ["Accessing VEA from the Windows Start menu"](#) on page 37.

See ["Accessing VEA through the Windows computer management window"](#) on page 37.

See ["Recommendations for caching-enabled disks"](#) on page 39.

See ["Disconnecting from a server"](#) on page 38.

Accessing VEA from the Windows Start menu

Use the following to access Veritas Enterprise Administrator (VEA) from the Windows Start menu and connect to a server.

To access VEA from the Windows Start menu and connect to a server

- 1 Launch VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

The Select Profile window appears.

- 2 Select a profile and click **OK** to continue.

The Veritas Enterprise Administrator window appears.

- 3 Click **Connect to a Host or Domain** to connect.

The Connect window appears.

- 4 Select a Host from the drop-down menu, click **Browse to find a Host**, or enter the Host name of the computer and click **Connect**.

You may also choose to connect using a different user account after selecting a Host. In this case, you are prompted for another user name and password to connect with.

Note: On the VEA GUI, in the **Connect** panel, if a period (.) is used instead of "localhost" as the **Host Name**, then the login credentials are not saved even if the **Save Password** option is selected while logging on.

Accessing VEA through the Windows computer management window

Use the following to access VEA through the computer management window and connect to a server.

To access VEA through the computer management window and connect to a server

- 1 Right-click the **My Computer** icon and select **Manage** from the context menu that appears.

The Computer Management window comes up.

The window's display can represent a local computer or a remote computer as follows:

- If the window represents the local computer, "(Local)" appears after the Computer Management node in the tree view.

- If the window represents a remote computer, the name of the remote computer appears after the Computer Management node.

The default setting is for the window to display the local computer as the top node.

- 2 To connect to a different computer, right-click the top node of the Computer Management tree view and select **Connect to another computer** from the context menu. In the **Select Computer** window that appears, click **Local** or **Another computer**.

For another computer, enter or browse for the name of the appropriate computer. Click **OK** to connect. You may be asked to enter a user name and password.

- 3 In the Computer Management tree view, expand the Storage node and the Volume Management node that appears under it.

If the server is running SFW or a version of Volume Manager for Windows, **Veritas Enterprise Administrator** appears under the Volume Management node.

- 4 Double-click **Veritas Enterprise Administrator**.

The Veritas Enterprise Administrator GUI comes up with the Select Profile window displayed.

- 5 Select a profile and click **OK** to continue.

The Veritas Enterprise Administrator console appears.

Disconnecting from a server

Use the following to disconnect from a server.

To disconnect from a server

- 1 Right-click the server icon you want to disconnect from, and select **Disconnect** from the computer context menu. You can also select **Disconnect** from the **File** menu or click the **Disconnect** tool on the toolbar, the second tool from the left (the small computer with a red checkmark).
- 2 A confirmation dialog appears; click **OK** to end the connection.

The computer icon disappears from the tree view.

If the connection is lost for any other reason, a message appears that indicates the connection has been terminated. SFW removes the icon for that computer from the tree view. You need to click **OK** to close the message box.

Understanding the conditions for remote connection

This section describes the conditions for connecting to a remote computer.

You can connect to a remote computer if it is established that your user name and password match those of an account with any of the following:

- Local Administrator rights of the remote server
- Domain Administrator rights of the domain of the remote server and you have logged on from that domain or a trusted domain
- Domain Administrator rights of a trusted domain to the remote server's domain and you have logged on from that domain or the remote server's domain

When you initiate the command to connect to a remote computer, Storage Foundation checks to see whether you have already met one of the three conditions above when you logged on initially to the local client computer. If you have, then you are connected directly to the remote computer without being asked to supply a user name and password.

You are not connected automatically if you have logged onto the local computer with an account that does not have administrator privileges on the remote machine you want to connect to. However, with the Connect command in the VEA GUI, you are always given the opportunity to enter a user name and password. Click the More button in the Connection dialog box to expand the dialog box and display the Authentication section. If you then enter a user name and password for an administrator account on the remote machine, you are connected to the remote computer.

Note: The first time you enter a user name and password, you can click the Remember password check box, so that you do not have to enter the password every time.

Recommendations for caching-enabled disks

The following describes how to disable write caches and HBA caches.

Write caches

Write caches for disks and disk arrays improves disk performance, but the cached data may be lost in an unplanned disk shutdown. In general, Veritas recommends that write caches be disabled to avoid any possible data loss. However, write caches may be enabled for disks and disk arrays that have an emergency backup power source for their caches.

Refer to your hardware documentation for additional information about disabling the write cache. Also refer to Microsoft knowledge base article #811392 (<http://support.microsoft.com/default.aspx/kb/811392>) for additional information to configure the write cache.

If you choose to have the write cache enabled, the warning message about write cache enabled disks is displayed during startup. This message can be turned off by clearing the check box for Show write cache enabled information in the **Storage Foundation General** tab of the Preferences GUI.

More information about turning off the message is available.

See “[The disk view](#)” on page 77.

It may be possible to disable the write cache for a disk by using the Windows Computer Management GUI as shown in the following steps.

To disable the write cache:

- 1 Select **Device Manager** in the left pane of Windows Computer Management.
- 2 Expand **Disk Drives** in the right pane to display the disks.
- 3 Right-click the appropriate disk and select **Properties** from the context menu.
- 4 Click the **Disk Properties** tab and clear the check box for **Write cache enabled**.
- 5 Click **OK** to save and close your setting.

HBA caches

Some storage devices, such as RAID controllers, may have HBA caches. If these storage devices are used in a cluster environment, such as Microsoft Failover Clustering or VCS, the data in the cache of a failed system has the potential of corrupting the data in the cluster. Veritas recommends that the HBA caches be disabled for these devices. Refer to your hardware documentation for information about disabling HBA caches.

Review the Veritas Enterprise Administrator GUI

Once you have connected to one or more servers, you can use the VEA GUI to view and manage the connected servers. This section gives a brief overview of the VEA GUI so that you may understand how to use the GUI to accomplish the tasks needed to set up and configure the Storage Foundation program on the server.

If you look in the right pane of the VEA GUI window, you should see an icon representing the server or servers that you are connected to. There are additional icons for all of the storage components associated with the server. By viewing these

icons, you can see at a glance how the different storage components on the server are organized and whether they function normally.

The key points about the VEA GUI are as follows:

- The tree view in the left pane provides an efficient overall view of the storage. In the VEA GUI, the different components of the storage being managed by Storage Foundation are represented as icons. The top level of the tree is the management console icon. Under it are the servers being managed by SFW. Under each server are the storage components associated with it—CD-ROMs, disk groups, disks, and volumes. By clicking on the + or - sign in front of an icon, you can display or hide subordinate icons under that icon. The Control Panel icon for each server gives access to further settings that can be applied to the storage objects. The Logs icon brings up an Event Log and a Task Log relating to the storage objects.
- If there is a problem with one of the storage components, a small error symbol is superimposed on the icon. A yellow caution symbol indicates there is a potential problem. A red circle with an x on it indicates that a serious storage failure has occurred, such as a disk failure.
See “[SFW error symbols](#)” on page 527.
- Selecting a storage object in the tree view brings up further information relating to that object in the right pane of the GUI window.
For example, if you expand the Disks icon in the tree view and click a particular disk, the right-pane view displays the volumes on that disk. That right-pane view for the disks has two other tabs, Disk Regions and Disk View. Disk Regions shows you subdisks (that is, parts of volumes that reside on the disk) and free space. The Disk View represents the same disk regions in a diagram and provides additional information. The tabs in the right pane change, depending on what storage object is selected in the tree view.
- Right-clicking a storage object in the tree view or right-pane view brings up a context-sensitive menu with the commands that are appropriate to that storage object.
For example, two tasks that are necessary in setting up storage under Storage Foundation are to create disk groups and volumes. You can right-click the Disk Group icon or a disk icon and a menu will come up that includes the command New Dynamic Disk Group. Likewise, once a dynamic disk group is created, you can click the Disk Group icon or an icon of a disk that is included in a dynamic disk group and a menu comes up that includes the command New Volume.
- You can also select commands from the menu bar and the toolbar.
Again, the commands on each menu can vary, depending on what storage object you have selected.

- The lower pane of the VEA GUI is used to display recent alerts relating to the storage objects, or it can also show the progress of storage tasks, such as formatting of a volume.
To toggle back and forth between the two purposes, click the Console or Tasks tab at the lower left corner of the pane.
More information about using the VEA GUI to manage your storage is available. See [“VEA Overview”](#) on page 65.
If you are installing on a new system with new disks or you have added new disks to the system, such disks show up in the tree view with the words "No Signature" after them. Right-click each disk and select Write Signature from the context menu.
See [“Add a disk signature to a disk”](#) on page 98.

Configure basic disks (Optional)

The Storage Foundation program defines disks as being either basic or dynamic. Basic disks are disks other than the ones that have been defined through SFW as dynamic. In most cases, it is not necessary to configure basic disks for your initial setup and configuration of SFW.

This section provides information about basic disks and describes the situations where you might want to use the commands provided within SFW to configure basic disks.

About basic disks

Basic disks are disks that have the standard partitioning scheme used by Windows NT, Windows 95/98, and MS-DOS. Disks configured on a fresh system are defined as basic disks. If the Windows operating system is installed, you have your system and boot partition already set up as a basic disk partition. This partition is created as a part of the Windows install process. Once Storage Foundation is running, you do not need to create any more basic disk partitions. You can then set up dynamic disk groups and volumes with your remaining storage to make them into dynamic disks. It is not necessary to place partitions on a new disk for it to be made into a dynamic disk through Storage Foundation.

Upgrading basic disks to dynamic

All basic disks can be upgraded to dynamic disks through Storage Foundation. To upgrade a basic disk to a dynamic disk, you need to create a dynamic disk group and include the disk in the group.

See [“About creating dynamic disk groups”](#) on page 43.

The basic disk does not have to include a partition on it; but if it does, a primary partition becomes a simple dynamic volume, and any logical drives within an extended partition also become simple volumes. The extended partition no longer exists.

The process of upgrading a basic partition to a dynamic volume is called "encapsulation," because it encapsulates the partition information within the dynamic volume structure.

You can also upgrade your system and boot partition to a dynamic disk and then mirror it.

See ["Set up a dynamic boot and system volume \(Optional\)"](#) on page 59.

About creating dynamic disk groups

When you create a dynamic disk group and add disks to it, those disks become dynamic—that is, they can now have Storage Foundation dynamic volumes created on them. Thus, a basic disk is said to be 'upgraded to a dynamic disk' through its addition to a dynamic disk group. You can change multiple basic disks to dynamic at one time when you use either the New Dynamic Disk Group command or the Add Disk to Dynamic Disk Group command.

See ["Considerations when changing a basic disk to dynamic"](#) on page 43.

See ["Creating a dynamic disk group"](#) on page 44.

See ["Making storage available after system startup"](#) on page 47.

See ["Add a disk to a dynamic disk group"](#) on page 99.

Considerations when changing a basic disk to dynamic

There are some considerations to make when you change a basic disk to dynamic.

Consider the following:

- Once you change a basic disk to a dynamic disk, the volumes on the disk cannot be accessed by MS-DOS, Windows 95/98, or Windows NT.
- After you upgrade a basic disk to a dynamic disk and create dynamic volumes on it, you cannot change the dynamic volumes back to partitions or basic volumes. You can revert a dynamic disk to basic, but you must delete all dynamic volumes on the disk and then use the Remove Disk from Dynamic Disk Group command.
- On Windows Server operating systems, you can upgrade a basic disk to a dynamic disk that can be contained in a Microsoft Disk Management Disk Group or in an SFW dynamic disk group as follows:

- If you upgrade a basic disk to a dynamic disk in a Microsoft Disk Management Disk Group, you can add it to an existing Microsoft Disk Management Disk Group or use it to create the Microsoft Disk Management Disk Group. You are not able to have more than one Microsoft Disk Management Disk Group.
- If you upgrade a basic disk to a dynamic disk in an SFW dynamic disk group, you can add it to an existing SFW dynamic disk group or use it to create an SFW dynamic disk group. You can have multiple SFW dynamic disk groups.
- To create a bootable dynamic disk that can be mirrored, you must upgrade a basic disk that contains the system partition to a dynamic disk. See [“Set up a dynamic boot and system volume \(Optional\)”](#) on page 59.
- When a basic disk with existing partitions is converted to a dynamic disk, all handles to the disk must be closed before the conversion. This means that all applications using the disk must be stopped before starting the conversion operation.

Creating a dynamic disk group

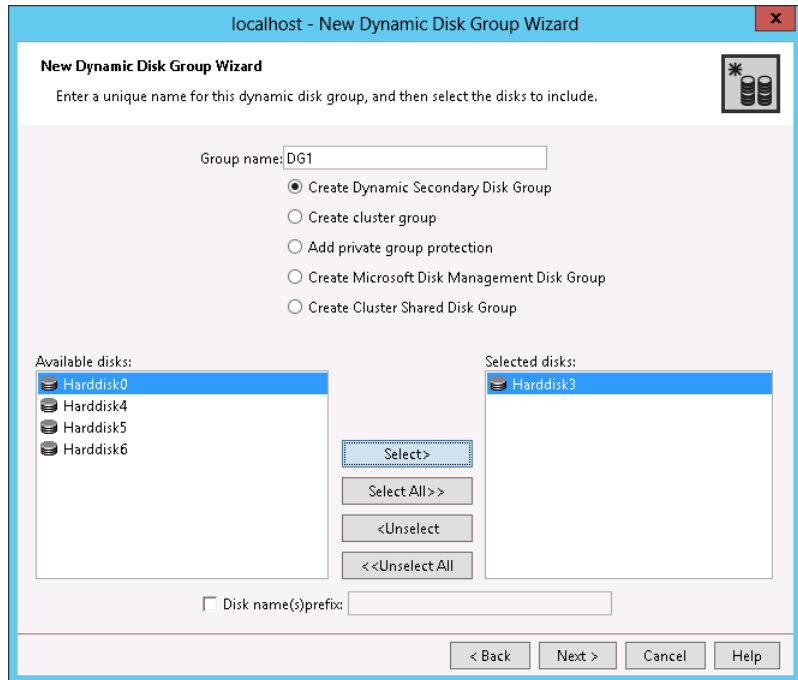
Use the following to create a dynamic disk group.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

To create a new dynamic disk group

- 1 In the VEA GUI toolbar, click **New Group**. Or, right-click the **Disk Groups** folder or a disk, and then select **New Dynamic Disk Group** from the context menu.
- 2 The New Dynamic Disk Group Wizard appears. Click **Next** to continue.

- 3 On the next panel, select the type of disk group that you want to create, provide a name for it and select disks for it.



Enter a name for the disk group, and indicate which disks you want to include in the group. Also, consider whether you want the disks to be part of a regular disk group or one of the optional disk group types, which are defined in the table below. If you click none of the optional disk group type checkboxes, you create a regular dynamic disk group.

Note: A dynamic disk group name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

Disk Group Type	Description
Create Dynamic Secondary Disk Group	Creates a dynamic (Secondary) disk group.

Disk Group Type	Description
Create cluster group	Creates a cluster dynamic disk group that can be used with VCS or Microsoft Failover Clustering. Note: If you create a disk group for configuring Microsoft Failover Clustering, be sure to check the Create cluster group check box.
Add private group protection	The feature of private dynamic disk group protection makes it possible to partition shared storage with a regular dynamic disk group.
Create Microsoft Disk Management Disk Group	For Windows Server operating systems, this creates a Windows Disk Management disk group.
Create Cluster Shared Disk Group	Creates a cluster-shared disk group (CSDG) that is used for Cluster Volume Manager (CVM).

Ensure that only the disks you want to include are listed in the **Selected disks** box.

Select the **Disk names prefix** check box and enter a disk name prefix to give the disks in the disk group a specific identifier. The identifier precedes a system-generated numeric ID that SFW uses to distinguish the disks in the disk group. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Click **Next** to continue.

- 4 The next panel confirms the disks that you have selected. Choose **Next** to continue. If required, you can click the **Back** button to go back to the previous panel to modify your disk choices.

For each disk, information is provided on whether it contains volumes and whether it is updated. If it is not going to be updated, you should modify the selection.

Normally, you would add all the disks you want in the group at this point. You can always add more disks later with the **Add Disk to Dynamic Disk Group** command.

- 5 On the next panel, verify the brief information about the disk group that is created and click **Finish** to create the disk group and exit the wizard.

Once a disk group is created, you can create dynamic volumes for the disk group.

Making storage available after system startup

Generally, dynamic (non-cluster) disk groups are available during system startup. However, some types of storage are not available until the service that controls the storage has been started. SFW provides the Veritas DG Delayed Import Service (VxDgDI) to accommodate this type of storage. VxDgDI and the `vxdg latestart` command allow SFW to access the storage when it becomes available.

See [“vxdg”](#) on page 593.

About creating dynamic volumes

This section on creating dynamic volumes has some useful background information, as well as the detailed steps for carrying out the procedure.

See [“Background information”](#) on page 47.

See [“Creating a dynamic volume”](#) on page 48.

See [“Dynamic volume types”](#) on page 56.

Background information

The following is background information to create dynamic volumes.

Dynamic volume layouts or types

There are five dynamic volume layouts or types—concatenated (simple or spanned), mirrored, striped, RAID-5, and mirrored striped (RAID 0+1).

See [“Dynamic volume types”](#) on page 56.

Volume size in the new volume wizard

It is important to understand that the total volume size you indicate for the volume in the New Volume wizard is the usable size for storing the data. Depending on the layout type, the volume can take more space on the disk. The size shown in the wizard for simple, striped, and spanned volumes is the actual volume size on the disk. A RAID-5 volume requires additional space for parity information, and a mirrored volume is a multiple of the space taken by the original plex of the volume.

The wizard provides a Max Size button to determine the size of the volume. Again, this is the usable size of the volume rather than the actual size taken by the volume on the disk.

Preselecting disks

You can preselect one or more disks before invoking the New Volume command. For example, you can click a specific disk in the tree view or select several disks in the right pane and then invoke the command from the context menu. If you have a disk or several disks selected, the program assumes that only the selected disks are used for the volume and the manual selection radio button will be indicated in the New Volume wizard. If you find that you want to use more disks than those originally selected, you can adjust the disk selection from within the wizard.

Thin Provisioned disks

Volumes can be created with disks that are enabled to support thin provisioning and storage reclamation. Using thin provisioned disks for a volume allows for efficient allocation of storage for the volume. A disk that supports thin provisioning is represented with a disk icon that includes a red colored sector. A disk that supports thin provisioning and storage reclamation is represented with a disk icon that includes a green colored sector with an asterisk (*).

Note: The operation to reclaim storage from thin provisioned disks is not supported for RAID-5, DCO, or Volume Replicator DCM volumes.

Note: The operation to reclaim storage is supported only for NTFS and ReFS volumes.

About dynamic volumes in SFW Basic

As in SFW, dynamic volumes can be created in SFW Basic. However, SFW Basic is limited to a maximum of four dynamic volumes. Upgrading SFW Basic to SFW lets you create any number of dynamic volumes.

Creating a dynamic volume

Use the following to create a dynamic volume.

To create a new volume

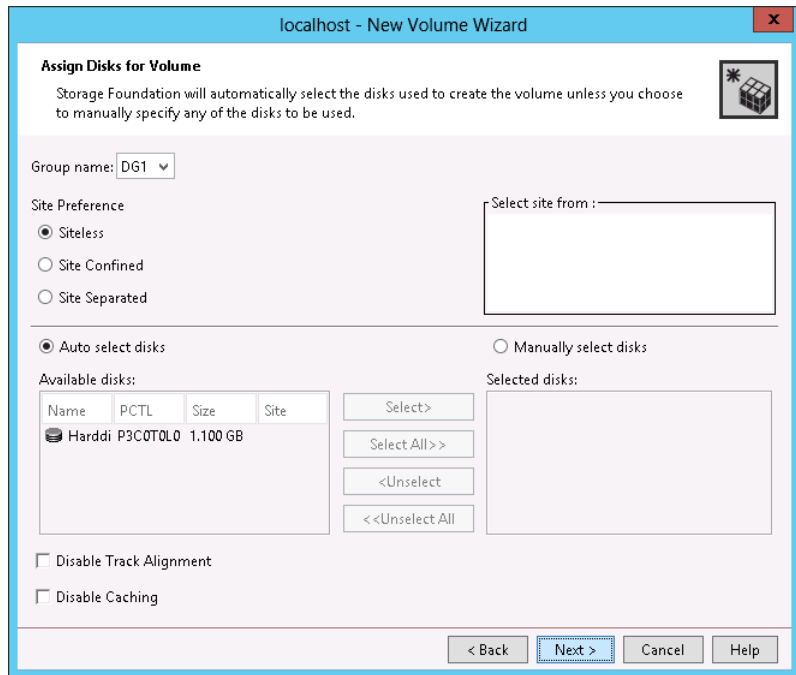
- 1** Right-click a dynamic disk in the tree view, the right pane, or the **Disk View** tab, and select **New Volume** from the context menu that comes up.

You can also select the command from the **Actions** menu or click the **New Volume** tool on the toolbar.

- 2** The New Volume Wizard appears.

Click **Next** to continue.

3 The Assign Disks for Volume panel appears. Complete the following:



Group Name A default dynamic disk group appears in the Group name box. If you want to create the volume in a different dynamic disk group, select it from the drop-down list.

Site Preference Specify the **Site Preferences** as **Siteless**, **Site Confined**, or **Site Separated** by enabling the required radio button against each of the respective options. For more information, See [“Administering volumes for site-based allocation ”](#) on page 456.

Select Site from Select the Site from the displayed list.

If you have specified Site Confined as Site Preference, then you can select at the most one site or maximum one site from the displayed list of sites.

If you have specified Site Separated as Site Preference, then you must specify at least two sites or more from the displayed list of sites. Press **Ctrl** on the keyboard to select multiple sites.

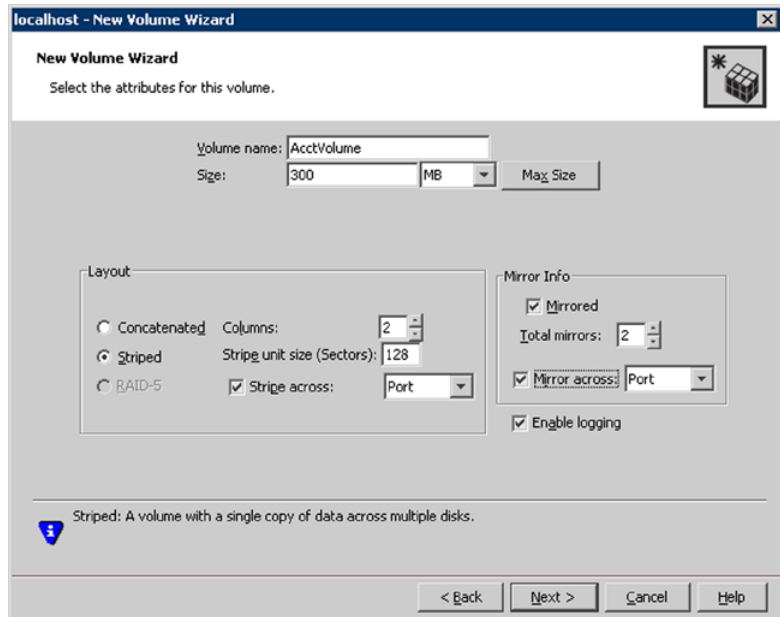
If a single site is selected for a Site Separated volume, then an error message to select at least 2 sites is displayed.

- Auto Select Disks** Automatically selecting disks is the default setting. Available disks are automatically assigned for volume creation.
- Manually select disks** Select this option to manually select disk or disks.
 If you had selected one or more disks when you invoked the New Volume command, the disk selection method is set to manual.
- Disable Track Alignment** You may also select **Disable Track Alignment** to disable track alignment for the volume.
 If you select this, then the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
 For more information about Track Alignment, See "[Track alignment](#)" on page 95.
- Disable Caching** SmartIO is a feature of SFW that helps in improving I/O performance of applications and Hyper-V virtual machines by providing a read-only I/O cache created on Solid-State Devices (SSDs).
 Caching is enabled by default for a new volume. To disable caching on the volume, select the **Disable Caching** check box.

Click **Next** to continue.

- 4** On this panel, select the attributes listed below for the volume.

Note: If you find the volume layout that you want is grayed out, it is most likely because not enough disks have been selected. Click the **Back** button and return to the previous panel to select more disks.



Volume Name

Enter a name for the volume in the **Volume name** text box.

The volume name is the internal Storage Foundation-specific name that is used in some SFW commands.

If you enter a volume name, SFW automatically assigns the file system volume label the same name. If you later change the file system volume label through the File System command, the Storage Foundation volume name and the file system volume label will be different. However, it will make little difference because the GUI identifies the volume by the file system label name except for the one entry that is listed as “Volume Name.” If you use the command line interface to create a volume, the file system name is not automatically set to the name of the volume.

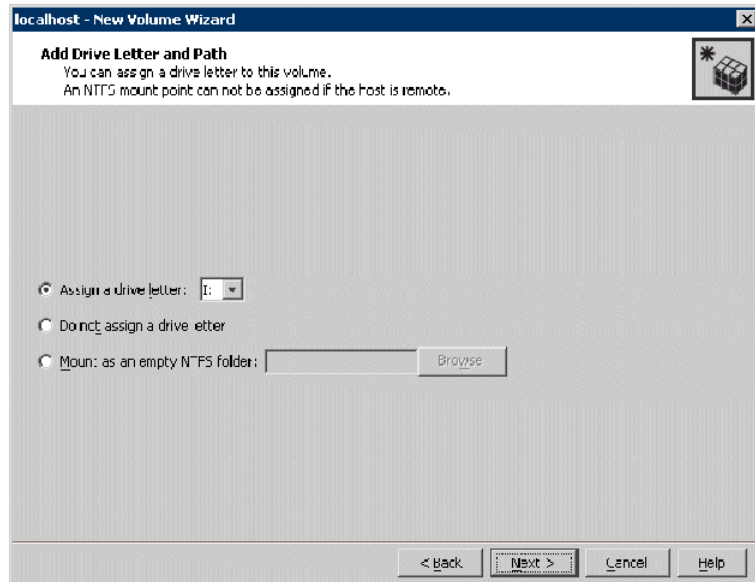
Note: A volume name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

Size	<p>Provide a size for the volume. You can use the pull-down list to the right of the Size entry box to select sectors KB, MB, GB, or TB for indicating the volume size.</p> <p>If you click the Max Size button, a volume size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.</p> <p>For a given dynamic disk group, the maximum volume size depends on the type of layout selected.</p>
Layout	<p>Select one of the following volume layouts:</p> <ul style="list-style-type: none">■ Concatenated If this option is selected for a Site Separated volume, then an error message is displayed prompting you to select a Mirrored layout. A Site Separated volume cannot be created without the mirror option.■ Striped■ RAID-5■ Mirrored Select this option for a Site Separated volume. A Site Separated volume cannot be created without the mirror option.■ Mirrored Striped (RAID 0+1) — choose Striped and the Mirrored check box If you create striped or RAID-5 volumes, the Columns and Stripe unit size must be specified. The stripe unit size is expressed as a number of sectors. Default values are provided. Additionally, you may specify to have the stripe across disks done by Port, Target, Enclosure, or Channel. <p>For a Concatenated or Striped volume, you may also specify a mirror volume with the following optional attributes:</p> <ul style="list-style-type: none">■ Number of mirrors■ Enable logging■ Mirror across disks by Port, Target, Enclosure, or Channel Note that the operation to create a new dynamic volume fails if the appropriate resources are not available to support the selected attributes to mirror across disks. See “Dynamic volume types” on page 56. <p>For more information about the volume layout types, See “Dynamic volume types” on page 56.</p>

Click **Next** to continue.

- 5 On the **Add Drive Letter and Path** panel, select one of the three choices described below.

Note: For cluster-shared volumes, ensure that the selected drive letter for the new cluster-shared volume is available and not in use on any of the cluster nodes.



- | | |
|-------------------------------|--|
| Assign a drive letter | Accept the default drive letter assignment or assign a drive letter by using the drop-down list. |
| Do not assign a drive letter | Select this option, if you may prefer to do this task later. |
| Mount as an empty NTFS folder | Select the option and provide a folder name or browse to select the folder name. This creates a new folder for you if you click the New Folder button in the Browse for Drive Path dialog box. |

Click **Next** to continue.

- 6 On the **Create File System** panel that appears next, indicate your choices for the file system selections that are described below.

Format this volume	<p>Select this check box if you want to format the volume using NTFS, FAT, FAT32, or ReFS.</p> <p>You can undo the selection if you want to format the volume later.</p>
Select a file system	<p>Make sure the file system type that you want is highlighted if you want to format the volume now.</p>
Allocation size	<p>If required, change the allocation size. The default size is recommended.</p> <p>Note: To set the appropriate allocation size for an application on your system, refer to the documentation for the application or Microsoft documentation concerning the application.</p>
File system label	<p>If required, type a file system label. If you do not enter a label, no default label is provided. You can enter a label for the file system later.</p> <p>Select any of the following options:</p> <ul style="list-style-type: none">■ Perform a quick format Select this option if you want to perform a quick format. If this check box is not selected, then the volume is formatted with normal formatting.■ Enable file and folder compression Select this option if you want to enable file and folder compression. This option is available only if the volume is formatted with NTFS. <p>Note: If you plan to use this volume to install SQL Server, do not select the Enable file and folder compression check box. The SQL Server installation cannot copy files on a compressed or encrypted folder.</p>

If you want to format the volume, create a file system label, or change any of the volume settings, right-click the volume and select **File System** and the subcommand **New File System** from the volume context menu. You can make your choices from a panel that has the same selections as the one above.

Click **Next** to continue.

7 Check your selections on the final panel and click **Finish**.

If you chose to format the volume, Storage Foundation begins the formatting process. Click the **Tasks** tab in the bottom left corner of the lower pane to see a text entry indicating the status and progress of the formatting process.

By default, the dynamic volume is created as a read/write volume. To make the dynamic volume a read-only volume, right-click the volume and check **Read only** in the Properties for the volume.

Dynamic volume types

This section provides a detailed description of the types of dynamic volumes that are available in Storage Foundation (SFW).

Note: In the SFW New Volume command, simple and spanned volume types are both referred to as concatenated and considered one type.

See [“Concatenated \(Simple and spanned\) volumes”](#) on page 56.

See [“Striped volumes”](#) on page 57.

See [“Mirrored volumes”](#) on page 57.

See [“RAID-5 volumes”](#) on page 58.

See [“Mirrored striped volumes \(RAID 0+1\)”](#) on page 59.

Concatenated (Simple and spanned) volumes

The concatenated volume type includes both simple and spanned dynamic volumes.

Simple volume	A dynamic simple volume consists of a single contiguous region (or subdisk) on a single physical disk. Dynamic simple volumes can be extended or mirrored. When you extend a simple volume to a noncontiguous region within the same disk or onto additional disks, it becomes a spanned volume.
---------------	--

Spanned volume A dynamic spanned volume consists of two or more subdisks (single contiguous regions) on one or more disks. With a spanned dynamic volume, you can combine sections of unallocated space from multiple dynamic disks into one large volume. The areas of unallocated space used to create spanned volumes can be different sizes. Spanned volumes are organized sequentially—that is, Storage Foundation sequentially allocates space on each disk until that disk is full and then continues with the next disk until the volume size is reached. Up to 256 disks can be included in a spanned volume.

You can expand the existing spanned volumes by the amount of unallocated space on all the disks in the dynamic disk group. However, after a spanned volume is extended, no portion of it can be deleted without deleting the entire spanned volume.

The advantage of a spanned volume is that it lets you create a large volume consisting of smaller pieces of disk space from several disks, thereby making more efficient use of the disk space than would be possible if you had to limit the volume to a single disk. The disadvantage of a spanned volume is that it is not fault tolerant. If one of the disks containing a spanned volume fails, the entire volume fails. However, a spanned volume can be mirrored.

Striped volumes

Striped volumes are created by combining areas of free space on two or more disks into one logical volume. Data is divided into blocks and spread in a fixed order among all the disks in the volume. With a striped volume, data is written to multiple disks, similar to spanned volumes. However, striping writes files across all disks so that data is added to all disks at the same rate. Up to 256 disks can be used in a striped volume.

Striped volumes offer the best performance of all the disk management strategies. However, as with spanned volumes, striped volumes do not provide fault tolerance. If a disk in a striped volume fails, the data in the entire volume is lost. However, a striped volume can be mirrored to provide data redundancy. In addition, Storage Foundation lets you extend a striped volume, either with or without a mirror.

Mirrored volumes

A mirrored volume is a fault-tolerant volume that duplicates your data on two or more physical disks. A mirror provides redundancy by simultaneously writing the same data onto two or more separate mirrors (or plexes) that reside on different disks. If one of the disks fails, data continues to be written to and read from the unaffected disk or disks.

Mirrored volumes protect the data on the disk from media failure. Mirroring reduces the chance of an unrecoverable error by providing duplicate sets of data, but it also multiplies the number of disks required for data storage and the input/output (I/O) operations when writing to the disk. However, some performance gains are achieved for reading data because of I/O load balancing of requests between the plexes. A mirrored volume is about the same as a RAID-5 volume in read operations but faster in write operations.

You can create a mirrored volume with the Create Volume wizard, or you can add a mirror to an existing volume with the Add Mirror wizard. To create a mirrored volume on two or more disks, each mirror or plex of the original volume requires space that is at least the same size as the original volume. When you create a mirror, the same drive letter is used for all plexes of the volume. If you have sufficient disks available, you can create multiple mirrors for extra redundancy (up to the limit of 32 mirrors). Then, you can break off a mirror to use for backup, data analysis, or testing without losing redundancy.

Note: Adding a mirror to a volume involves a certain amount of time for mirror resynchronization. The SmartMove feature can help reduce mirror resynchronization time.

See [“SmartMove”](#) on page 93.

Breaking off a plex of the mirrored volume does not delete the information, but it does mean that the plex that is broken off no longer mirrors information from the other plex or plexes in the mirrored volume. The broken-off plex is assigned a different drive letter than that of the original mirrored volume. It is no longer part of the mirrored volume, but it retains its other volume layout characteristics. For example, if you had a mirrored striped volume, the broken-off plex would become a striped volume.

In the case of an unrecoverable error on a plex within a mirrored volume, you need to remove the plex that is damaged with the Remove Mirror command. You then can use the Add Mirror command to create a new plex on another disk to take the place of the damaged plex.

When you want to use the space in a mirrored volume for other purposes, you can remove a plex and return the space that it used to unallocated space.

Storage Foundation lets you add a mirror to any kind of dynamic volume, except RAID-5. You can extend any mirrored volume.

RAID-5 volumes

A RAID-5 volume is a fault-tolerant volume with data and parity striped alternately across three or more physical disks. If one of the physical disks fails, data that was

on the failed disk is recreated from the remaining data and parity on the other disks. RAID-5 volumes are a good solution for data redundancy in a computer environment in which most activity consists of reading data.

You need a minimum of three disks for a RAID-5 volume. Storage Foundation can accommodate up to 256 disks in a RAID-5 volume. You can extend a RAID-5 volume, but it cannot be mirrored.

Mirrored striped volumes (RAID 0+1)

RAID 0+1 volumes are mirrors of striped volumes. For example, a two-disk striped volume can be mirrored to two additional disks. This RAID type provides the advantages of both speed (from striping) and fault tolerance (from mirroring). You can add more mirrors to a mirrored striped volume, and you can extend this type of volume onto additional dynamic disks within the dynamic disk group.

Set up a dynamic boot and system volume (Optional)

You may want to make the Windows boot and system partition into a dynamic volume and mirror that volume. Then, if the disk containing the boot and system volume fails, you can start the computer from the disk containing the mirrors of the volume.

To make a boot and system partition dynamic, you include the disk that contains the basic active boot and system partition in a dynamic disk group. When you do that, the boot and system partition is automatically upgraded to a dynamic simple volume that is active—that is, the system will boot from that volume.

For an optimal setup, it is recommended that you have a separate disk for the dynamic system and boot volume and mirror it with one or two disks. Thus, you will have a boot dynamic disk group with two or three disks. You then need to have an additional disk or disks for the data. The data disks would be in another dynamic disk group.

In Windows documentation, the boot volume is the partition that contains the operating system and the system volume is the partition that the computer starts from. The boot and system volumes can be in the same partition or different partitions for MBR style partitions. For GPT style partitions, the boot and system volumes must be in different partitions.

This following procedure is written for an MBR style partition with the boot and system volumes in the same partition. If the boot and system volumes are in different partitions, then these steps can be used as a guide to apply to each volume separately.

For GPT style partitions, the following procedure can also be used as a guide to apply to each volume separately. However, use the Microsoft bootcfg.exe utility instead of editing the boot.ini as described in the procedure.

Note: SFW does not support a dynamic system volume on an Itanium (IA64) system with a GPT style partition. Therefore creating a mirror of a dynamic system volume on an IA64 system with this configuration is not supported. Refer to the *Veritas InfoScale Release Notes* for more information about creating and booting from a mirror on a system with this configuration.

Note: On Windows Server operating systems, although you can create and manage the dynamic boot and system volume using the VEA GUI, the native Windows driver handles the input/output operations for that volume. Therefore, ensure that you comply with Windows Server OS information on best practices and restrictions on the use of a dynamic boot and system volume.

Setting up a dynamic boot and system volume for MBR style partitions:

- 1** Make sure that the server has a basic system and boot partition that is active.
- 2** As a precaution, create a boot floppy of the system partition.

If something goes wrong with the original system and boot volume or its mirror, you can boot your system from the floppy. Test the boot floppy while your system is in a known good state.
- 3** Make sure that there are sufficient disks attached to the server so that there is a separate disk for the boot and system volume plus one or two more for mirroring the system disks and then additional disks for the data.

- 4 Create a dynamic disk group that has the disk with the existing boot and system partition plus one or two other disks to be used for mirroring purposes.

See [“About creating dynamic disk groups”](#) on page 43.

When you designate a disk as part of a dynamic disk group, the entire disk becomes a dynamic disk—that is, a disk capable of having dynamic volumes. Any basic partitions on that disk become dynamic volumes. Thus, the boot and system partition automatically becomes a dynamic simple volume. You do not have to mark it as the active volume, because it is automatically made active.

There is no command in Storage Foundation for making an existing dynamic volume active—that is, to make it the volume that the computer starts from. The only way to make a dynamic volume active through SFW is to upgrade the existing active basic system partition by including the disk that contains the partition as a member of a dynamic disk group. You can make a basic partition active through the command **Mark Partition Active**.

- 5 If you have not done so already, create one or more additional disk groups for your data and then create the necessary volumes for the data.

See [“About creating dynamic disk groups”](#) on page 43.

See [“About creating dynamic volumes”](#) on page 47.

- 6 Create one or more mirrors on the dynamic system and boot volume.

See [“Add a mirror to a volume”](#) on page 208.

The troubleshooting section covers the difficulties that can occur.

See [“An attempt to mirror a boot or system volume fails or has an error message”](#) on page 549.

- 7 If the dynamic boot or system volume fails, you must make the change to the server’s **boot.ini** file to make one of the mirrored boot or system volumes the active volume for your computer and then restart.

If a break-mirror operation is performed on a mirrored boot volume, the resulting new volume—the broken-off mirror—is not usable as a boot volume.

Set desired preferences

This section describes some of the settings in SFW. Defaults are indicated where appropriate. Some settings require input before they are enabled.

See [“Hot relocation settings”](#) on page 62.

See [“Event notification settings”](#) on page 62.

See [“S.M.A.R.T. monitoring for a disk”](#) on page 63.

See [“Capacity monitoring for volumes”](#) on page 63.

See [“Task throttling”](#) on page 63.

See [“Configuration task performance tuning”](#) on page 64.

See [“Track alignment settings”](#) on page 64.

Hot relocation settings

Hot relocation is an automatic process that relocates subdisks from a disk that has I/O errors. If there are I/O errors on a subdisk containing a redundant volume (RAID-5 or mirrored), the redundant volume’s subdisks are moved from the failed subdisk to a hot spare disk. If there is not enough free space on designated hot spare disks, then free space on any available disk is used.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Hot Relocation.

Note: Dynamic disks belonging to a SAN Disk Group do not support Hot Relocation.

The Undo Hot Relocation command can be used to move the subdisks back to their original location after the failed disk has been repaired or replaced.

More information about how hot relocation works is available.

See [“Hot relocation”](#) on page 216.

The default for Storage Foundation is to have automatic hot relocation mode inactive. This means that if an I/O error occurs in a redundant subdisk, the subdisk is not automatically relocated to another disk. Use the Set Disk Usage command to designate preferred disks as targets for hot relocation.

The option to disable the hot relocation mode is available from the Control Panel.

See [“Hot relocation mode”](#) on page 218.

Event notification settings

SFW provides event notification by SMTP email, by pager, and through the SNMP traps that can be displayed in HP OpenView, CA Unicenter, and IBM Tivoli. You can configure the notification service to send messages to specific individuals, to groups, or to a management console in the case of SNMP traps. The event notification service is implemented through SFW’s Rule Management utility. If you want to set up event notification, you must use the Rule Management utility to set

up the rules that send out notifications after certain events occur. You access the Rule Management utility through SFW's Control Panel.

See ["Event notification"](#) on page 190.

The details of defining the rules for events are in the online Help in the module "Optional Rules for Handling Events." To invoke the online Help, choose Contents from the Help menu in SFW. In the tree view that comes up in the left pane, you see the help module "Optional Rules for Handling Events."

S.M.A.R.T. monitoring for a disk

Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) is available on many new SCSI disks. S.M.A.R.T. provides a way for disks to self-check certain physical conditions that are useful predictors of disk failure and send an alert to the SFW Event Log when conditions indicate that a failure may occur soon. This allows the system administrator to take preventive action to protect the data on the disk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support S.M.A.R.T. monitoring.

To receive alerts in the SFW Event Log, S.M.A.R.T. monitoring must be enabled for the global configuration and also on each disk you want checked. The default is to have S.M.A.R.T. monitoring disabled.

See ["Disk monitoring"](#) on page 193.

Capacity monitoring for volumes

Capacity monitoring is a feature that monitors the capacity of SFW dynamic volumes so that when used space on any volume reaches certain size thresholds, you are notified with an alert message. The purpose of capacity monitoring is to provide you with a warning when any dynamic volume has nearly reached full capacity. You can set up capacity monitoring for individual volumes or all volumes. Capacity monitoring is turned off by default. You can set individual volumes to grow automatically when a certain threshold is met.

See ["Capacity monitoring"](#) on page 195.

Task throttling

This feature lets you control how the system allocates CPU time between I/O and other tasks that must also be performed.

See ["Task throttling"](#) on page 94.

Configuration task performance tuning

Configuration task performance tuning lets you set the number of I/O threads to perform certain mirror-related operations, such as creating a mirror or adding a mirror.

See [“Configuration task performance tuning”](#) on page 95.

Track alignment settings

This feature lets you set dynamic volumes to always store blocks of data in alignment with the boundaries of the physical track of the disk. Aligning data so that it does not straddle a track boundary helps optimize I/O performance.

You can specify track alignment through the VEA GUI.

See [“Track alignment”](#) on page 95.

You can also specify track alignment through the CLI.

See [“vxdisk”](#) on page 624.

Protecting your SFW configuration with vxcb

To protect your SFW configuration (that is, the disk group and volume layouts), it is highly recommended that you perform periodic backups of the configuration, using the `vxcb` utility.

More information is available about configuration backup and `vxcb`.

See [“Configuration backup”](#) on page 88.

See [“vxcb”](#) on page 708.

Using the GUI to manage your storage

This chapter includes the following topics:

- [VEA Overview](#)
- [Systems perspective](#)
- [The disk view](#)
- [Search](#)
- [Assistant perspective](#)
- [Logs perspective](#)
- [Configuring SFW](#)

VEA Overview

The Graphical User Interface (GUI) for Storage Foundation for Windows is known as the Veritas Enterprise Administrator (VEA).

The VEA console window provides a graphical way to view and manipulate all the storage objects in your system. The VEA GUI display for Storage Foundation for Windows storage objects is shown in the sample screen that follows.

The VEA console display has the following features:

- Perspectives of the system
- A tree view displaying the storage objects in the left pane of the window
- Tabbed views in the right pane of the window containing additional information on the storage objects

- A menu bar, toolbar, and URL bar
- A lower pane containing information on processes that the program initiates or a listing of alert messages

Viewing disks and volumes

The VEA console is used to view disks and volumes and their relationship with each other.

Viewing all disks associated with a volume

All the disks associated with a volume can be accomplished with the following steps.

To view all disks associated with a volume

- 1 In the left pane, click the desired volume icon.

If the volume icon is not displayed because the tree view is not fully expanded, click the plus sign in front of the nodes starting at the top of the tree. The Volumes icon appears under the StorageAgent node.
- 2 If necessary, in the right pane of the console window, select the **Disks** tab.

The **Disks** tab view displays the disk or disks associated with a particular volume.

Viewing all volumes associated with a disk

All the volumes associated with a disk can be accomplished with the following steps.

To view all volumes associated with a disk

- 1 In the left pane, click the desired disk.

If the disk icon is not displayed because the tree view is not fully expanded, click the plus sign in front of the nodes starting at the top of the tree. The Disks icon appears under the StorageAgent node.
- 2 If necessary, in the right pane of the console window, select the **Volumes** tab.

The **Volumes** tab view displays all the volumes that are on that disk.

SFW general preferences

SFW General preferences are settings for displaying GUI features relating to SFW. To access these settings select Preferences from the VEA Tools menu. In the dialog box that appears, select the Storage Foundation General tab, as shown in the screen below.

Check the desired settings and click **Apply** to enable the settings in the VEA GUI. Click **OK** to exit the dialog box. To make your selected settings appear in the VEA GUI, you may need to refresh the right pane. Click an object in the left pane to refresh the right pane.

The available settings are as follows:

Show inline disk view	Allows the Disk View to be accessed as a tab in the right panel of the VEA GUI.
Show dialog top information panel	Enables the display of helpful information at the top of certain dialogs.
Launch task dialogs as modal	Task dialogs are launched so that the task dialog is the only active window on the desktop.
Enable subdisk drag and drop	Allows the drag and drop of subdisks to a different location on the same disk or to a different disk in the Disk View.
Show write cache enabled information	If SFW determines there is a write cache enabled disk, then this setting enables the display of a message reminding the user about write cache enabled disks every time SFW is launched.

The URL bar

The URL bar, located above the tabs of the right pane, displays the selected object's location in the tree in a URL style. If you are connected to more than one host, you can click on the drop-down to select a different host.

The Toolbar

The Toolbar provides quick access to the program's most common tasks.

Table 3-1 Description of Toolbar Icons







Icon	Name	Description
	Connect	Brings up the Connection dialog box.
	Disconnect	Disconnects from selected machine.
	New Window	Opens a new window that displays the console.

Table 3-1 Description of Toolbar Icons (*continued*)

Icon	Name	Description
	New Volume	Brings up the New Volume wizard.
	New Group	Brings up the New Dynamic Disk Group wizard.
	Search	Searches by specific criteria for volumes, disks, or disk groups.

The Perspective bar

The Perspective bar located at the far left of the console, provides quick access to different perspectives (views) of the system that you are connected to. Clicking a perspective displays certain operations and objects of the system.

The perspectives are as follows:

System	Displays the objects, details of the objects, and status of the system.
Assistant	Provides shortcuts to perform common tasks on the objects of the system. This provides a task-based approach to operations that can be performed on the system.
Logs	<p>Displays the system's task and alert logs in the right pane. Clicking the Alert Log tab or the Task Log tab in the right pane displays the corresponding log entries.</p> <p>The logs are as follows:</p> <ul style="list-style-type: none"> ■ The Task Log is a listing of tasks, such as formatting a volume or resynchronization of a volume. In SFW, the Task Log is not implemented. ■ The Alert Log lists events and alerts pertaining to the Storage Foundation for Windows program. <p>See “Alert log” on page 189.</p>
Control Panel	<p>Displays the tasks related to configuring the system. The Control Panel lets you access settings for Logs, Rule Management, Configuration Backup, Historical Statistics, SMTP Configuration, Task Settings, and Track Alignment.</p> <p>See “Configuring SFW” on page 87.</p>

Systems perspective

The Systems perspective displays the SFW console and it consists of the following:

- See “The left pane” on page 69.
- See “The right pane” on page 72.
- See “The lower pane” on page 76.

The left pane

In the System perspective, the left pane shows a tree view of the system and storage objects detected by the Storage Foundation for Windows software. The tree view displays the hierarchical relationships of the objects. The node at the top of the tree represents the Storage Foundation for Windows client that you are connected to. In the screen below, the client is connected to "localhost." The objects under this node are the managed servers that the client is connected to and managing. In the screen below, there is only one managed server node, a server named "jktestmachine."

Below each managed server icon are the following object categories:

- | | |
|--|-----------------|
| Default | ■ Cache |
| | ■ CD-ROMs |
| | ■ Disk groups |
| | ■ Disks |
| | ■ Enclosures |
| | ■ Saved Queries |
| | ■ Volumes |
| Systems configured for support of Microsoft multipath input/output (Microsoft MPIO) solution | ■ DMP DSMs |
| Systems running VSS-aware applications, such as Microsoft SQL Server | ■ Applications |

The tree view can be expanded by clicking on a plus sign (+) in front of an object icon. When the tree view is fully expanded, all the objects have a minus (-) sign in front of them. By clicking on a minus sign at any level, you can collapse an object down to that level. The fully collapsed tree shows only the top-level object.

Right-clicking on an object in the tree view brings up a context menu that is appropriate to that object.

The following is additional information about the storage object categories under each managed server node.

Cache	Cache area is the storage space allocated on the SSD(s) for caching. It is used to store cache data corresponding to any caching-enabled volume.
CD-ROMs	Any CD-ROM drives recognized by Storage Foundation for Windows as existing on the computer you are managing.
Disk groups	A disk group is a grouping of disks within Storage Foundation for Windows. The two types of disk groups are basic and dynamic. See “Disk groups overview” on page 146.
Disks	Disks are physical disks or logical disks recognized by the Windows operating system. Depending on the type of disk, a disk may be enabled to support thin provisioning and storage reclamation. Thin provisioning is a technology to efficiently allocate storage for a disk. Thin provisioning allocates physical storage only when actual data is written to the disk. Some disks that are enabled for thin provisioning also provide storage reclamation. Storage reclamation is the operation that decreases the physical storage allocation once data is deleted from the disk. A disk that supports thin provisioning is represented with a disk icon that includes a red colored sector. A disk that supports thin provisioning and storage reclamation is represented with a disk icon that includes a green colored sector with an asterisk (*).
Enclosures	Enclosures are physical objects that contain one or more physical disks. For example, the disks may be contained in arrays or JBODs. Also the disks may be internal to your server.
Saved Queries	Saved Queries refers to queries that were saved with the Search feature of SFW. If you saved queries with the Search feature, then this node would display the results of the saved query. See “Search” on page 80.

Volumes	<p>A volume is a logical entity that is made up of portions of one or more physical disks. A volume can be formatted with a file system and can be accessed by a drive letter or a mount point. Storage Foundation for Windows works with basic and dynamic volumes.</p> <p>A volume may be either read only or read/write. The icons for read-only volumes include a picture of a padlock to differentiate them from read/write volumes. Not all commands available in Storage Foundation for Windows for read/write volumes are enabled for read-only volumes because specific commands require write access to the volume. Check the access mode of a particular volume if a command is not available.</p>
DMP DSMs	<p>On the servers that are configured for support of Microsoft multipath input/output (Microsoft MPIO) solution, a node for DMP DSMs appears.</p> <p>Completely expanding the DMP DSMs node displays DSM nodes being used, nodes of arrays being controlled by the DSM, and the disks contained in the array. These nodes let you manage the settings for the arrays and disks configured for Microsoft MPIO.</p> <p>See “DMP overview” on page 399.</p>
Applications	<p>On the servers that are running VSS-aware applications, such as Microsoft SQL Server, a node for Applications appears.</p> <p>SFW provides an option of taking snapshots with Volume Shadow Copy Service (VSS). The VSS snapshot method lets you take snapshots of VSS-aware applications, such as Microsoft SQL Server, while the application files are open. When VSS-aware applications do not exist, the snapshot is taken with the SFW FlashSnap method (VM method).</p> <p>See “Volume Shadow Copy Service (VSS)” on page 289.</p> <p>See “Using the Enterprise Vault Snapshot wizard” on page 327.</p> <p>See “Using the VSS snapshot wizards with Microsoft Exchange” on page 295.</p> <p>See “Using the VSS snapshot wizards with Microsoft SQL” on page 342.</p>

- iSCSI
- On the servers that are connected to an iSCSI SAN, the following nodes may appear:
- iSCSI Initiators
This node displays information about the initiators and targets of the iSCSI SAN.
 - iSCSI Arrays
This node displays information about the Microsoft iSCSI Target subsystems of the iSCSI SAN.
 - iSNS
This node displays information about the iSNS Server of the iSCSI SAN.
- See [“Overview of iSCSI SAN support”](#) on page 176.

The right pane

In the System perspective, the right pane provides information on the various storage objects. The tabs available in this pane are determined by the object you have selected in the tree view in the left pane. In each tab view, you can right-click to get a context menu of commands available in that view. In every view, you can widen or narrow columns by placing the mouse pointer on the border between two columns in the header row and moving the mouse to the left or right while holding down the mouse button.

This section focuses on how information about disks, volumes, and cache area managed by Storage Foundation for Windows displays in the right pane.

See [“Right-pane display for disks and volumes”](#) on page 72.

See [“Disk view tab”](#) on page 74.

See [“Right-pane display for cache”](#) on page 75.

See [“Right-pane legends”](#) on page 75.

See [“Customizing the table headers”](#) on page 76.

Right-pane display for disks and volumes

This section describes how the tabbed views change, depending on the disk and volume storage objects you select.

Note: Information about thin provisioned disks is not automatically updated in the VEA GUI. The information about thin provisioned disks can be updated by performing a Refresh operation.

Selecting all disks

In the tree view in the left pane of the VEA console, if you select all disks by clicking the Disks folder, the right-pane view displays tabbed views for Disks and Disk View. The Disks tab shows information about the disks available on the selected server, as shown in the sample screen below. The Disk View tab uses a graphical format to show information about the volumes contained on each disk. The Alerts tab displays information about alerts.

Selecting an individual disk

If you select an individual disk in the tree view, you see a display similar to the sample screen below. Notice that the tabs have now changed to Volumes, Disk Regions, Disk View, and Alerts. The Volumes tab gives information about the volumes on the selected disk.

Viewing subdisks

If you click the Disk Regions tab, you see all the subdisks or parts of a volume that reside on the disk, as well as any free regions on that disk. In the screen below, there are eight subdisks.

Every volume has at least one subdisk. A striped volume has a minimum of two subdisks. A RAID-5 volume has a minimum of three subdisks. A mirrored volume has as many subdisks as there are mirrors, which are also called "plexes."

Note: There are no subdisk storage objects in the left-pane tree view. You must access the subdisks through the right pane. You have to click an individual disk in the left pane to gain access to the **Disk Regions** tab.

Selecting volumes

If you select all volumes by highlighting the Volumes folder in the tree view, the right pane displays the tabs: Volumes, Disk View, and Alerts. When you select an individual dynamic volume in the tree view, you see a right-pane view similar to the screen below. Notice how selecting an individual volume makes available many more tabbed views.

A dynamic volume has the following tabbed views:

Table 3-2 Table of Volume Tabbed Views

Tab	Description
Disks	Displays the information about the disks in the volume.

Table 3-2 Table of Volume Tabbed Views (*continued*)

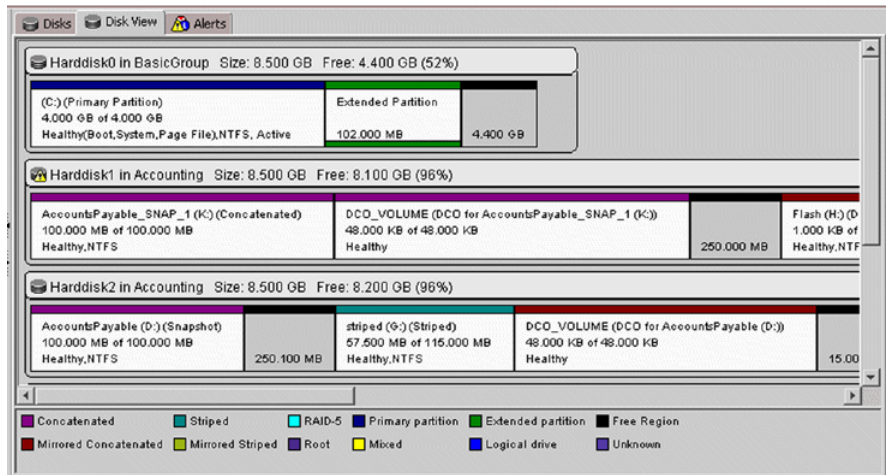
Tab	Description
Mirrors	Details the mirrors in the volume. See “Mirroring” on page 206.
Logs	Shows any logs that are added to a volume to increase the speed of resynchronizing mirrored or RAID-5 volumes after disk failures. See “Dirty region logging (DRL) and RAID-5 logging” on page 222.
Subdisks	Lists the subdisks associated with the volume. Subdisks are the basic units in which SFW allocates disk space in a volume. Each subdisk represents a specific portion of a volume on a dynamic disk. For example, a mirrored volume with two mirrors (plexes) has two subdisks.
Snapshot Volumes	Provides information on snapshots. Snapshots are mirrored volumes that have been broken off from their original volume for backup or other purposes. This tab appears only when you have created snapshot volumes. See “Snapshot commands” on page 259.
DCO	Provides information on the Disk Change Object (DCO) volumes that are used with FastResync and also with FlashSnap. See “Disk Change Object (DCO) volume” on page 256.
Disk View	A detailed view of the disks associated with the volume. See “Disk view tab” on page 74.
Alerts	Displays information about alerts. This tab appears only when an event has occurred that generates an alert.

A volume that is not dynamic, such as a partition on your C drive, displays only the Disks, Disk View, and Alert tabs.

Disk view tab

The Disk View tab displays a graphical layout of the disks on your system, including removable media. As with the tree view, right-clicking on an object or portion of the window relating to an object brings up a context menu that has commands related to the object.

The Disk View is shown below.

Figure 3-1 Disk view tab

The Disk View is also available as a separate window.

See [“The disk view”](#) on page 77.

Right-pane display for cache

To view details of an individual cache area, expand the Cache node and select the desired cache area.

If you select an individual cache area in the tree view in the left-pane of the VEA console, the right-pane view displays the tabbed views for Disks and Volumes. The Disks tab displays the details of the SSD devices that are used to create the cache area. The Volumes tab displays the details of the volumes linked with the selected cache area.

Right-pane legends

This section describes the two legends in the right pane, the disk usage legend and the Disk View legend.

Disk usage legend

The disk usage legend, that is located at the bottom of the right pane, is context-sensitive and changes depending on the tab selected in the right pane. The legend shows either disk usage or the volume type.

When the Disks, Subdisks, or Disk Regions tab is selected, the disk usage legend shows how to interpret symbols showing the level of I/O activity (Normal, High, or Critical) in each area.

Disk view legend

The legend shown in the Disk View or with the Disk View tab selected shows the volume or partition type. The colors represent the following layout characteristics: concatenated, striped, RAID-5, primary partition, extended partition, free region, mirrored concatenated, mirrored striped, root, mixed, logical drive, and unknown.

The mixed type refers to a volume that has both concatenated and striped plexes because it has been modified by using the Dynamic Relayout feature.

See [“Dynamic relayout”](#) on page 225.

Customizing the table headers

The Customize Columns window lets you customize the table headers by changing which columns are shown and in what order.

To customize the table headers

- 1 Right-click the column header.
- 2 Select **Customize Table Header**. The Customizing Column Settings window appears.

Headers of columns shown in the right pane of the console are shown on the right side of the dialog. The available headers are shown on the left side.

- 3 Make the desired changes to the column headers.

The options are as follows:

- To add a column header, select it in the **Available Columns** pane and click the **Add** button.
- To remove a column header, select it in the **Selected Columns** pane and click the **Remove** button.
- To rearrange the sequence of the headers, select the header you want to move in the **Selected Columns** pane and use the **Up** and **Down** buttons to move it within the list.

- 4 Click the **OK** button to apply the settings. The column headers displayed in the console are in the order you selected in the **Selected Columns** pane.

The lower pane

In this section of the System perspective, you can toggle the display by clicking on either the Console tab or the Tasks tab at the bottom left corner. The default Console tab display shows Storage Foundation for Windows alerts. The Tasks tab shows progress on Storage Foundation for Windows tasks, such as formatting a volume.

The name of the selected managed server is displayed in the bottom center. This is helpful if you have connected to multiple servers with SFW.

The disk view

This section describes the Disk View.

Disk view overview

The Disk View can be accessed in two ways.

It can be accessed as follows:

- Select the Disk View tab in the right pane.
- Right-click an object (disk group, disk, or volume) in the tree view pane, and select Disk View from the context menu that appears.

The same commands are available in both views. The Disk View available from the tree view context menu opens a separate window that can remain open as you continue to work in Storage Foundation for Windows.

Additionally, you can use the Ctrl key to select only the specific objects (disks, volumes, or disk groups) you want to include in the separate disk view. This feature is helpful if you have a very large number of disks or volumes.

The commands in Disk View include Refresh View, Print, Show Legend, and four options for different views of the selected disks: Expand, Vol Details, Full Volume Display, and Projection. These commands are available from several different places.

From the Disk View tab in the right pane, right-click in an empty space (right-click any area of the pane where there is not a graphical representation of a disk), and a context menu containing the commands appears.

If you have created a separate window by selecting Disk View from the context menu, these commands can be accessed by right-clicking in an empty space (right-clicking on any area of the pane where there is not a graphical representation of a disk), by clicking the Options or File menu on the menu bar, or by clicking on icons on the toolbar.

Refresh view

The Refresh View command in the **Disk View** menu updates the Disk View only and ensures that any recent changes are represented accurately.

Show legend

The legend is located at the bottom of the Disk View. The Show Legend command toggles the Disk View legend on and off. If you work with a large number of disks, you may want to turn the legend off to gain additional space in the window. Each volume type or partition type is assigned a color. The legend shows the color-coding scheme. When the Vol Details command is selected, there is a colored bar across the top of each subdisk that indicates its volume type.

The disk view legend is shown below.

Figure 3-2 Disk view legend



Print

The Print command sends your current Disk View display to a selected printer. It may be convenient to have a hard copy of your disk configuration, especially if you work with a large number of disks.

The remaining items on the **Disk View context** menu: Expand, Vol Details, Full Volume Display, and Projection are discussed in the sections that follow.

Options for graphical views

There are four options for different views of the selected disks: Expand, Vol Details, Projection, and Full Volume Display. The table below shows the options available for these commands and for the print function. The toolbar containing these icons is displayed only in the separate Disk View window. However, in both versions of the Disk View, you can right-click an empty space (an area of the pane where there is not a graphical representation of a disk) to bring up a context menu with these options.

Table 3-3 Disk View Options





Icon	Meaning	Tasks
	Represents the Expand function in the Disk View.	Lets you view detailed information about disks and the organization of the data on the disks. If this icon is not checked, the view is contracted to show less details.

Table 3-3 Disk View Options (*continued*)

Icon	Meaning	Tasks
	Represents the Vol Details function in the Disk View.	Displays subdisks as partial volumes.
	Represents the Projection function in the Disk View.	Highlights objects associated with the selected subdisk or volume.
	Represents the Print function in the Disk View.	Lets you print the contents of the current view.

Note: The Full Volume Display view, which displays more detailed information about subdisks, is not represented by an icon.

Expanded view

The subdisks can be viewed in either the expanded view or nonexpanded view. Each view gives different information. When Expand is checked in the Disk View menu, the expanded view is in effect. When it is not checked, the nonexpanded view displays.

Expanded view

The expanded view shows the subdisk names and sizes, as well as the amount of free space remaining on the disk.

Nonexpanded view

If the expanded view is not selected, a contracted disk view is shown. This nonexpanded view lets you easily see the number of subdisks on each disk.

Vol details view

The Vol Details view adds information to either the expanded or nonexpanded view. This view is useful if you look at a large number of disks, because each disk takes less space than in the full volume display, so more disks can be displayed on the screen at once.

In the nonexpanded view, Vol Details shows a colored bar at the top of each subdisk that indicates its volume type. The legend at the bottom of the pane shows the color-coding scheme. The legend can be toggled on and off with the Show Legend command.

Additionally, in the expanded view Vol Details shows the volume name and drive letter associated with each subdisk.

Full volume display

As shown below, Full Volume Display provides even more information about the subdisks. It is available only in the expanded view and only when the Vol Details option is selected. In addition to the volume name and type, Full Volume Display shows the volume size and status.

Projection

Projection highlights all the subdisks associated with a selected volume. After the Projection command is selected, click a subdisk. The selected subdisk becomes dark purple, and the other subdisks in that volume become a lighter purple. In the example below, the projection feature was turned on, and the subdisk named "Mirrored (F:)" on Harddisk3 was selected. That subdisk became dark purple, and the other mirror that is part of the volume "Mirrored (F:)" (located on Harddisk4) became a lighter purple. This feature can be useful for viewing a volume that is spread across several disks.

Context menu available from disk view

The context-sensitive menu of commands that is available by right-clicking on a volume or disk in the tree view is also available by right-clicking a subdisk or disk in the Disk View. Additionally, in the separate Disk View window, the context-sensitive commands are available on the **Action** menu.

Also, in Disk View you can drag and drop subdisks instead of using the Subdisk Move command (if drag and drop is enabled in the general preferences).

Search

The Search feature lets you find specific volumes, disks, or disk groups on a large system or to find a list of volumes, disks, or disk groups meeting specified criteria.

Note: Searches are case-sensitive.

Conducting a search

The following describes the search procedure.

To conduct a search

- 1 Click **Search** on the VEA console toolbar or select **Tools > Search** from the menu bar to open the Search window.
- 2 Click the **New Query** tab to display the available search criteria.
- 3 Select the object type from the drop-down list to the right of the words "Look for:" (near the top left corner of the window). The choices are Volumes, Disks, or Disk Groups.
- 4 Select the attribute you want to search for. The available attributes depend on the object type.

Attributes for Volume

Name

Volume Name

Disk Group Name

Status

Device

FS Capacity

FS Label

FS Free Space

Attributes for Disk

Name

Disk Group Name

Status

Size

Free Space

Attributes for Disk Group

Name

Imported

Size

Free Space

Note: "FS" means "File System."

5 Define the limiter you want to use.

The following options limit the searches:

Any	Finds all instances of the selected attribute. It is not necessary to enter more criteria in the text box below the limiter box. For example, searching Disk Groups for any Free Space returns all disk groups with free space available.
Contains	Finds any of the selected attributes that contain the string that is entered in the text box below the limiter box. Contains acts as a wildcard (*) and searches for the string anywhere within a longer string. For example, searching Volumes for a name that contains "Stripe" returns all volumes that have the string "Stripe" anywhere in the volume name, including "Striped (H:)" and "JKStriped (J:)." Because searches are case-sensitive, "Stripe" does not find "A striped volume (L:)."
Is	Finds only an exact match with the string entered in the text box below the limiter box. It is useful for finding a specific volume, disk, or disk group on a very large system. For example, searching Disks for "Name is Harddisk1" returns Harddisk1. You must enter the exact information, including the drive letter or location for volumes.
Is not	Excludes the specific information entered in the text box below the limiter box. For example, searching Volumes for "FS Type is not FAT32" returns all volumes that are not formatted as FAT32.

6 Enter the search string in the text box below the limiter box. Searches are not case-sensitive.

7 Click **Add** to create the search term, which appears in the box to the right.

If more than one search term is defined, they are joined by a logical "and" operator; that is, the result contains all of the search terms. For example, a search for the terms "Volumes with Free Space" and "any NTFS Volumes" results in a list of all volumes that both are formatted as NTFS and contain free space.

The logical "or" operator is not provided in this search function.

The **Remove** button can be used to delete a highlighted search term.

The **Save Query** button lets you save the query you created for future use.

Use the following steps to save a query:

- Save query
- Click **Save Query** to open the Save Query window.
 - Enter a name to identify the query you want to save.
 - Enter a brief description of the search that the query performs.
 - Check the **Active** check box if you want the results of the query displayed in the VEA console.
 - Click **OK** to save the query.

- 8** Click **Search Now** to complete the search.

The search results appear in the lower pane.

You can double-click a specific storage object (disk, volume, or disk group) in the search results to bring up the information in the VEA console.

- 9** Click **Save Results** to save the results of the search.

Do the following in the **Save** dialog box:

- Save dialog
- Select the folder where you want to save the file.
 - Enter a name for the file in the **File name** field.
 - Click **Save** to save the file and close the dialog box. The file is saved as a text file with a .txt extension.

- 10** Click **Clear All** to delete all the search terms and results. This allows for a new search.

Using saved queries

The following procedure describes the way to use saved queries.

To use saved queries:

- 1** On the VEA console toolbar, click **Search**, or select **Tools > Search** from the menu bar to open the Search window.
- 2** Click the **Saved Queries** tab in the Search window.
- 3** Select a previously saved query from the drop-down list to the right of the words "Query:" (near the top left corner of the window). The choices are all the previously saved queries.

If a description of the query was saved with the query, it is displayed in the Description message box.

- 4 Click **Search Now** to run the search.

The search results appear in the lower pane.

You can double-click a specific storage object (disk, volume, or disk group) in the search results to bring up the information in the VEA console.

- 5 To save the results of the search, click **Save Results**.

Perform the following steps in the Save dialog:

- Save results
- Select the folder you want to save the file to.
 - Enter a name for the file into the **File name** field.
 - Click **Save** to save the file and close the dialog box. The file is saved as a text file with a .txt extension.

- 6 Click **Clear All** to delete the search results and run another search.

Creating search reports

This section describes how to create reports from the results of the Search feature. Using the Search command, it is possible to generate a wide variety of reports that provide information on the disk groups, disks, and volumes on each server. The results of these searches can be saved and imported into Microsoft Excel or any other spreadsheet or database application that accepts tab-delimited text files. The results can be sorted, displayed as graphs or charts, and imported into written reports or PowerPoint presentations.

Search reports provide a snapshot of a system at a particular point in time. By running the same search on a regular schedule (monthly, annually, etc.), information can be collected, evaluated, and used to better manage the system.

The following sample report was created by using the Search feature and Microsoft Excel.

Determine the categories you want to include in your report, and use the Search feature to output the results to a text file.

This search lists the free space available on all the volumes located on the accessed server. After the search results are saved by using the Save Results button, a Microsoft Excel report can be generated.

Follow these steps to generate the search report:

- 1 Open the Excel application.
- 2 Click File > Open to open the .txt file saved by the SFW search.

3 You have to select All Files (*.*) from the Files of type drop-down menu. Navigate to the correct file and click **Open**.

4 Do the following in the Text Import wizard that appears:

- Text Import Wizard
- Select the file type that best describes your data.
 - Select Tab to be the delimiter.
 - Select the General data format.
 - Click **Finish** to complete your report in Excel.

Refer to the Microsoft Excel help for more information about the commands and features available in Excel.

Types of search reports

The type of search and search report you generate depends on the object type you base your search on.

Table 3-4 Attributes Available for Volume, Disk, or Disk Group Search

Attributes for a Volume Search	Attributes for a Disk Search	Attributes for a Disk Group Search
Name	Name	Name
Volume Name	Disk Group Name	Imported
Disk Group Name	Status	Size
Status	Size	Free Space
Device	Free Space	
FS Capacity		
FS Label		
FS Free Space		

Note: "FS" means "File System."

Each search is run only on the specific server selected. For convenience, you can write a macro to automate the Excel steps in the report. By opening the search text files from multiple servers in Excel and running the macro, you can generate a combined report for a group of servers. The table that follows lists some sample searches and the types of report that would be generated.

Table 3-5 Sample Searches and Results

Search	Criteria	Report Shows
Volume	FS Free Space, any	All volumes that contain free space
Volume	Status is degraded	Volumes whose state is degraded and that might fail soon
Disk	Size greater than 10 GB	All disks larger than 10 GB
Disk	Free Space greater than 1 GB	All disks with more than 1 GB of free space
Disk	Disk Group is DynGrp1	All disks in DynGrp1
Disk Group	Name, any	All disk groups. Useful to run the search on multiple servers

Assistant perspective

The purpose of the Assistant is to make available the most commonly used commands in the program.

To access the Assistant, select Assistant in the Perspective Bar. Then select the object from the drop-down located on the left side of the window that you want to perform tasks on. To perform tasks on disks, disk groups, and volumes select StorageAgent from the drop-down.

Command icons appear in the pane located on the right side of the window when StorageAgent is selected. Clicking a command icon launches the dialog for the command.

The type of commands displayed in the Assistant is controlled by selecting an object in the object pane and an activity in the activity pane located on the left side of the window. For example, selecting Disk in the object pane and Modify in the activity pane displays commands that pertain to modifying disks. Selecting All Objects in the object pane and selecting All Activities in the activity pane would display all the commands that are available in the Assistant.

Logs perspective

The Logs perspective displays the Alert Log and the Task Log messages associated with the SFW storage objects. Select the appropriate tab to access the log you are interested in.

See [“Alert log”](#) on page 189.

Configuring SFW

Use the following to configure Storage Foundation.

Control Panel perspective

To access the Control Panel, select Control Panel in the Perspective bar. Expanding the tree view displays the system objects. Selecting one of the system objects displays the Control Panel objects for that system object in the right pane.

Like the Windows Control Panel, this Control Panel has icons that open up into windows with settings for various functions within the program.

Selecting localhost displays the following:

- See [“Log settings”](#) on page 88.

Selecting actionagent displays the following:

- See [“Rule Manager”](#) on page 88.

Selecting StorageAgent displays the following:

- See [“Configuration backup”](#) on page 88.
- See [“Historical statistics”](#) on page 89.
- See [“SMTP configuration”](#) on page 89.
- See [“System settings”](#) on page 89.

System settings include the following:

- See [“Hot relocation”](#) on page 91.
- See [“Snap Shot”](#) on page 91.
- See [“GUID recycle ”](#) on page 92.
- See [“SCSI Support”](#) on page 92.
- See [“SmartMove”](#) on page 93.
- See [“Disk Management Snap-in”](#) on page 93.
- See [“Task settings”](#) on page 94.
Task settings include the following:
 - See [“Task throttling”](#) on page 94.
 - See [“Configuration task performance tuning”](#) on page 95.
- See [“Track alignment”](#) on page 95.

To access any of these functions, double-click the appropriate icon in the Control Panel.

Log settings

Double-clicking the Log Settings icon brings up the Log Settings dialog box. This dialog box lets you set the maximum amount of storage used for the Alert Log and the Task Log.

The Alert Log lists events or alerts pertaining to the Storage Foundation for Windows program.

The Task Log is a listing of tasks, such as formatting a volume or resynchronizing a volume. The Task Log Settings window lets you set the maximum log size and to specify a policy when the maximum log size is reached.

See [“Log configuration”](#) on page 189.

Rule Manager

Double-clicking this icon brings up the Rule Manager window, that lets you configure rules for event notification messages and actions.

See [“Event notification”](#) on page 190.

Configuration backup

Double-clicking the Configuration Backup icon brings up the Configuration Backup window. Configuration Backup can preserve the disk group, dynamic disk, and volume configuration information that can be used in the recovery of a disk.

More information about recovering a disk is available.

See [“vxcbr”](#) on page 708.

Configuration Backup only backs up the configuration of disk groups that are online. This operation does not process any offline or deported disk groups.

Enter the following settings in the **Configuration Backup** window to back up the configuration of disk groups:

- Automatic check box

The following explains the Automatic check box:

Check the check box	Back up operation runs every time that a configuration change occurs.
---------------------	---

Clear the check box	Back up operation only runs when you click OK.
---------------------	--

- **Number of copies** The number of backup copies of the configuration to create. (Range of copies: 5-1024.)
- **Number of minutes between backups**
The interval of time between backup copies, when Automatic is selected. (Range of minutes: 60 minutes - 10080 minutes (1 week).) This setting is ignored when Automatic is not selected.
- **Directory for automatic backup**
The path of the target directory for the backup copies, when Automatic is selected.
- **Directory for manual backup**
The path of the target directory for the backup copies, when Automatic is not selected.

Click **OK** when all the settings have been entered. (Clicking **Reset** before clicking **OK** recalls the previously saved settings.)

Historical statistics

Double-clicking the Historical Statistics icon brings up the **Historical Statistics Settings** dialog box. This dialog box lets you define the file size for the files that are used for historical data collection. It also has a setting for the data collection rate.

See [“Configuring settings in the historical statistics settings window”](#) on page 238.

SMTP configuration

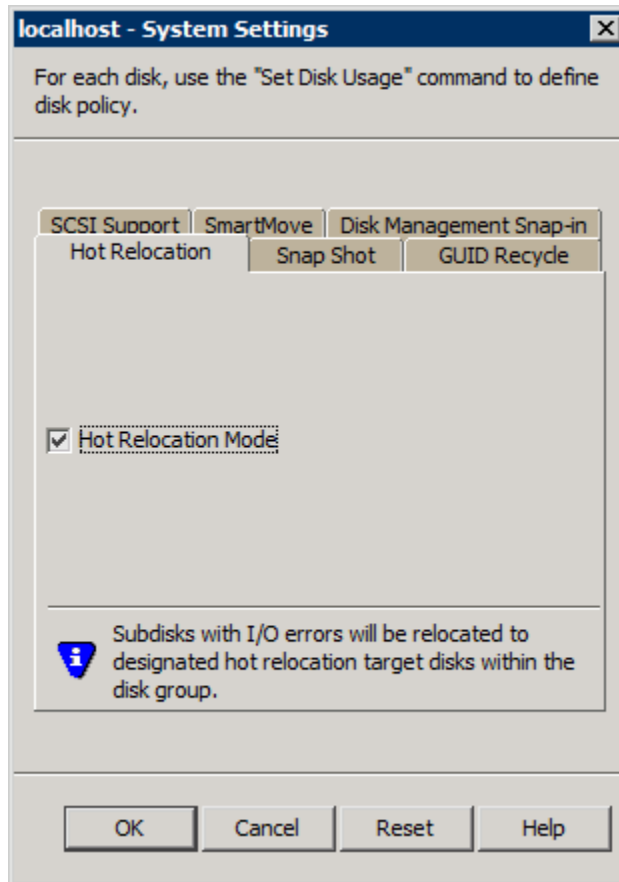
Double-clicking the SMTP Configuration icon brings up the SMTP Settings dialog box. This dialog box lets you define the IP address of the SMTP mail server on your network so that email notifications can be sent automatically. It also has settings for the port of the SMTP mail server and the email address that is used for sending email.

See [“SMTP configuration for email notification”](#) on page 204.

System settings

Double-clicking the System Settings icon brings up the System Settings dialog box. The dialog box lets you enable the hot relocation volume feature, set the volume snapshot method, set GUID recycling, set SCSI support, and enable SmartMove.

Figure 3-3 System Settings dialog box



System Settings tab

Description

Hot Relocation	To enable the Hot Relocation feature.
Snap Shot	To select which snapshot method to use for taking snapshots.
GUID Recycle	To enable GUID recycling.
SCSI Support	To enable SCSI support.
SmartMove	To enable the SmartMove setting.

System Settings tab

Disk Management Snap-in

Description

To replace the Disk Management Snap-in in the Windows Computer Management and Server Manager consoles with the SFW VEA GUI.

Hot relocation

Clicking the Hot Relocation tab brings up the Hot Relocation window. Check the Hot Relocation Mode check box to enable the feature. When a subdisk containing a redundant volume fails, hot relocation automatically moves the redundant volume from the failed disk to hot spare disks, or to other free space if no space is available on hot spare disks.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Hot Relocation.

Note: Dynamic disks belonging to a SAN Disk Group do not support Hot Relocation.

See [“Hot relocation”](#) on page 216.

Snap Shot

Clicking the Snap Shot tab brings up the Snap Shot window. Click **VSS** or **VM** to select the snapshot method to use for taking snapshots.

VSS and VM snapshot methods offer the following:

- The VSS snapshot method lets you create copies of volumes for VSS-aware applications, such as Microsoft SQL Server. For VSS-aware applications, the VSS snapshot method quiesces the application before taking the snapshot. For applications that are not VSS-aware, the VSS snapshot method does not quiesce the application before taking the snapshot. The VSS snapshot only creates read-only snapshots.
- The VM snapshot method lets you create copies of volumes for any application. The copies may be either read only or read/write depending on the setting of the volume being copied. In other words, a snapshot of a read-only volume results in a read-only snapshot and a snapshot of a read/write volume results in a read/write snapshot.

GUID recycle

Clicking the GUID Recycle tab brings up the GUID Recycle window.

A Globally Unique Identifier (GUID) is a Windows convention to identify components and objects in the system. In some environments, the number of GUIDs in the system may exhaust the storage in the Windows registry. Enabling volume GUID recycling helps to avoid this situation by using GUIDs from a GUID recycle list. The GUID recycle list is a list of GUIDs that have been previously used, but are no longer needed by the system.

To enable GUID recycling, select Enable Volume GUID Recycling and the desired settings as follows:

GUID recycling statistics	Enables the monitoring of how many times a GUID from the GUID recycle list is used.
Check for duplicate GUIDs	Enables a check to see if a GUID selected from the GUID recycle list is already in use in the system.
Automatically fix duplicate GUIDs	Enables the generation of another GUID when a GUID is found to be already in use in the system.

SCSI Support

Clicking **SCSI Support** brings up the SCSI Support window.

Click **SCSI-2 support** or **SCSI-3 support** to select the protocol setting to support your storage devices in a clustering environment (If you select the Microsoft Failover Cluster option during the product installation, SCSI-3 is selected by default).

Selecting SCSI-3 enables SFW clustering support to issue SCSI-3 commands to your SCSI-3 enabled storage devices. For this selection, setting the DMP DSMs for SCSI-3 support has no effect.

More information about DMP DSMs is available.

See [“DMP overview”](#) on page 399.

Selecting SCSI-2 enables SFW clustering support to issue SCSI-2 commands to your storage devices. If your storage devices are enabled for SCSI-3 commands, the DMP DSMs translate the SCSI-2 commands into SCSI-3 commands, if the DMP DSMs setting for SCSI-3 support is enabled.

Note: The SCSI Support window in the SFW Control Panel does not enable/disable SCSI-3 support of DMP DSMs.

More information about enabling SFW SCSI support for Microsoft Failover Clustering is available.

See [“Additional considerations for SFW Microsoft Failover Clustering support”](#) on page 503.

SmartMove

Clicking the SmartMove tab brings up the SmartMove window.

Check **Enable SmartMove** to enable the SmartMove setting to help reduce the resynchronization time for mirror operations and subdisk moves. SmartMove reduces the resynchronization time by using the NTFS or ReFS file system metadata to resynchronize only those regions that the NTFS or ReFS file systems use. When SmartMove is not enabled, SFW resynchronizes all regions.

Note: SmartMove supports only NTFS and ReFS disks.

Disk Management Snap-in

Clicking the Disk Management Snap-in tab brings up the Disk Management Snap-in window.

Check Replace Disk Management Snap-in to replace the Disk Management Snap-in in the Windows Computer Management and Server Manager consoles with the SFW VEA GUI. The SFW VEA GUI becomes accessible in both the Windows Computer Management console and the Windows Server Manager console.

Uncheck Replace Disk Management Snap-in to access the Disk Management Snap-in in the Windows Computer Management and Server Manager consoles.

Also, in a similar way, the following commands can be used to make the Disk Management Snap-in or the SFW VEA GUI accessible in the Windows Computer Management and Server Manager consoles.

Replace Disk Management with VEA GUI	<code>regsvr32 "%VMPATH%\extens.dll"</code>
Enable Disk Management	<code>regsvr32 /u "%VMPATH%\extens.dll"</code>
Replace Disk Management with VEA GUI (silent mode)	<code>regsvr32 /s "%VMPATH%\extens.dll"</code>
Enable Disk Management (silent mode)	<code>regsvr32 /s /u "%VMPATH%\extens.dll"</code>

Note: Silent mode suppresses a Windows pop-up after the command is entered.

Task settings

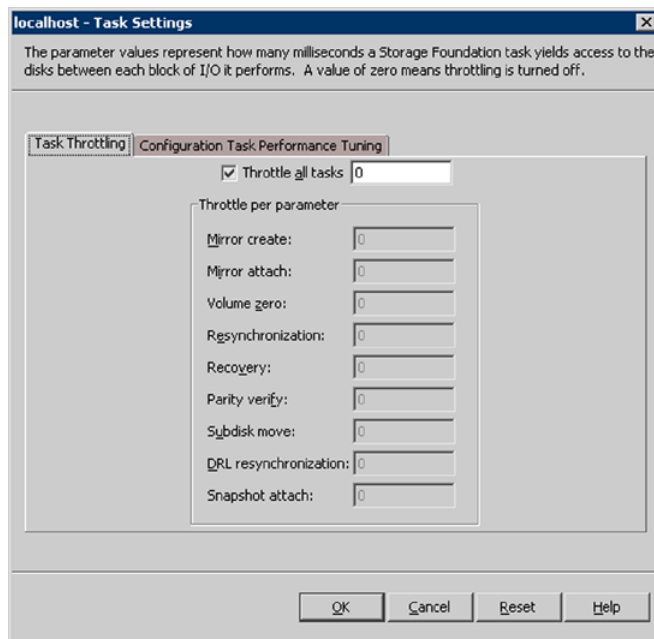
Double-clicking the Task Settings icon brings up the Task Settings dialog box. The dialog box lets you enable the task throttling feature and set the configuration task performance tuning feature.

Task throttling

Clicking the Task Throttling tab brings up the Task Throttling window. Task throttling lets you determine the priority of certain tasks.

The Task throttling tab is shown below.

Figure 3-4 Task Throttling tab



Using task throttling causes an operation to insert pauses for the specified amount of time during I/O. These pauses help to reduce any interference that may occur with the I/O of other applications running on the system. By selecting the Throttle all tasks check box, you apply the time delay that is specified in the text field to all Storage Foundation tasks. To apply the time delay to individual tasks, uncheck the check box, enter the number of milliseconds in each task's text field, and click OK.

Depending on your total system I/O activity, a time delay of 300 to 500 milliseconds may result in a favorable balance of I/O for Storage Foundation tasks and other applications. The Reset button restores default settings.

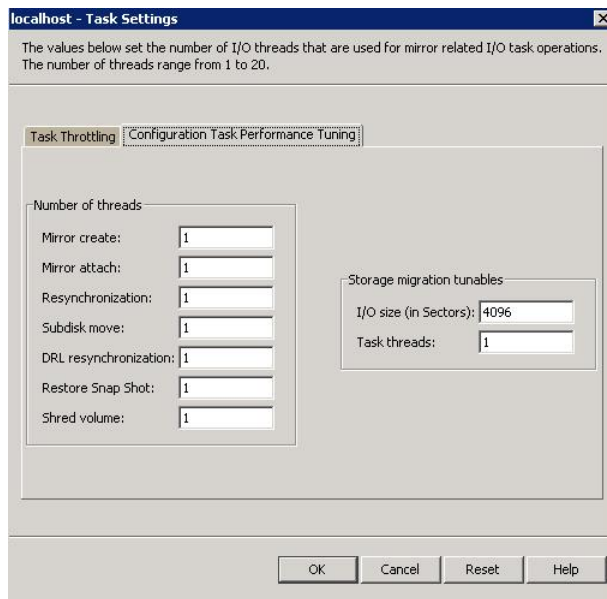
See [“Configuration task performance tuning”](#) on page 95.

Configuration task performance tuning

Clicking the Configuration Task Performance Tuning tab brings up the Configuration Task Performance Tuning window.

The Configuration Task Performance Tuning tab is shown below.

Figure 3-5 Configuration Task Performance Tuning tab



Configuration task performance tuning lets you set the number of I/O threads to perform certain mirror-related operations, such as creating a mirror or adding a mirror. The number of threads specified affects the performance of the operation.

Storage migration tunables let you optimize the storage migration performance:

See [“About performance tunables for storage migration”](#) on page 478.

Track alignment

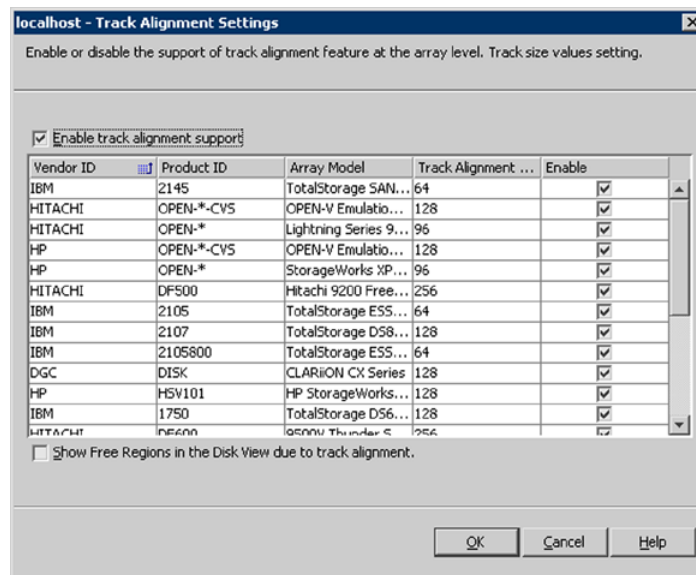
Double-clicking the Track Alignment icon brings up the Track Alignment Settings dialog box.

This feature lets you set dynamic volumes to store blocks of data in alignment with the boundaries of the physical track of the disk. Storing data so that it is aligned with track boundaries helps optimize I/O performance. Check the Enable track alignment support check box to enable the feature for the arrays listed.

You may also check the Show Free Region check box to display additional region information in the Disk View GUI.

The Track alignment settings are shown below.

Figure 3-6 Track alignment settings



Arrays not explicitly listed can be enabled as Default. (SFW sets all arrays in the Default category a vendor ID of 'default' and a product ID of 'default'.) This setting may or may not optimize I/O performance for the array. Contact your array manufacturer for details about track alignment and offset values.

You can use the CLI command, `vxdisk set track` to adjust the setting for track alignment for arrays enabled as Default.

See [“vxdisk”](#) on page 624.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Working with disks, partitions, and volumes

This chapter includes the following topics:

- [Overview](#)
- [Adding storage](#)
- [Disk tasks](#)
- [General Partition/Volume tasks](#)
- [Basic disk and volume tasks](#)
- [Automatic discovery of SSD devices and manual classification as SSD](#)
- [Volume Manager space allocation is SSD aware](#)

Overview

This chapter describes how to manage disks, partitions, and volumes.

Adding storage

This section describes the commands that are used while adding storage to your system.

See [“Add a new disk to your system”](#) on page 98.

See [“Add a disk signature to a disk”](#) on page 98.

See [“Add a disk to a dynamic disk group”](#) on page 99.

See [“Create new volumes”](#) on page 101.

Add a new disk to your system

Perform the following steps to add a new disk.

To add a new disk to your system

- 1 Install the new physical disk or disks.
- 2 From the **Actions** menu, select **Rescan**.

See [“Update disk information by using rescan”](#) on page 104.

On Windows Server operating systems, you also need to bring the disk online to make the disk accessible. To bring the basic disk online, right-click the disk to bring up the context menu and select the **Online** command.

- 3 After the operation is finished, the VEA GUI displays current information.

If Storage Foundation for Windows does not detect the new disk, you may need to restart.

A brand new disk has no signature or partition style. You can add a signature and partition style by using the **Write Signature** command.

More information about partition styles is available.

See [“What can you do with a basic disk?”](#) on page 134.

New disks are added to the computer as basic disks. You can upgrade new disks to dynamic disks at any time.

Add a disk signature to a disk

A new disk must have a signature on it before it can be used.

When you install a new disk and then do a Rescan, the new disk appears under the Disks icon. You can see the disk with the words "(No Signature)" in the tree view. If left unsigned, the disk cannot be used.

Once you bring the disk online and a signature is added to it, the disk displays as a basic disk. You can create partitions on the basic disk, or you can upgrade the disk to dynamic to let you create dynamic volumes on it.

For earlier version of Volume Manager, the command Upgrade a Dynamic Disk was used to change a basic disk to a dynamic disk. For later versions of the program, use the command Add Disk to Dynamic disk group to change a basic disk to a dynamic disk.

See [“Add a disk to a dynamic disk group”](#) on page 99.

You can also make a basic disk into a dynamic disk by including it as a member disk when creating a new disk group with the New Dynamic disk group command.

See [“About creating dynamic disk groups”](#) on page 43.

To add a signature to a disk

- 1 To add a signature, first you need to bring the disk online. To do this, right-click the unsigned disk, and then select **Online Disk**.
- 2 Right-click the disk again, and then select **Write Signature** from the context menu.

The Write Disk Signature dialog box appears. Select the disks you want to add a signature to.

To select a disk, click on it in the list of available disks and click **Add**. To select all the disks, click **Add All**.

To remove a disk from the **Selected disks** list, select the disk in the **Selected disks** list and click **Remove**. To remove all the disks from the **Selected disks** list, click **Remove All**.

- 3 Select the partition style of the disk by clicking the radio button for one of the following:
 - MBR (Master Boot Record)
 - GPT (GUID Partition Table)

For more information on both partition styles, refer to the Microsoft documentation.

- 4 Once you have selected the disks and style, click **OK**. Click **Cancel** to exit without adding a signature to the disk or disks selected.

Add a disk to a dynamic disk group

If you want to add more disks to a dynamic disk group after the group is created, use the Add Disk to Dynamic disk group command. You can also use this command to first create a new dynamic disk group and then specify the disks to be added to it.

To add a disk to a dynamic disk group

- 1 Right-click a basic disk in the tree view or the **Disk View** tab in the right pane, and then select **Add Disk to Dynamic Disk Group**.
- 2 The Add Disk to Dynamic disk group wizard appears. Click **Next** to continue.

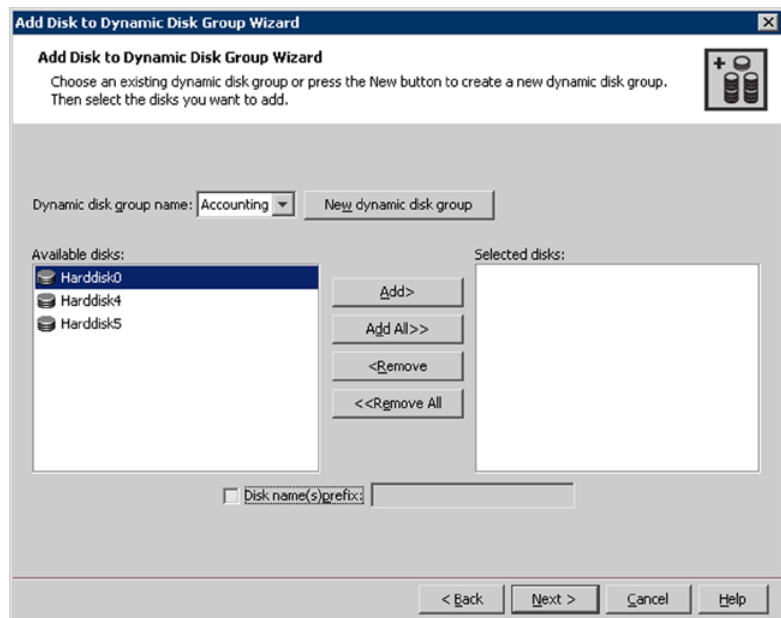
- 3 In the next screen, select one or more disks and the dynamic disk group you want to add them to.

Optionally, check the **Disk names prefix** check box and enter a disk name prefix to give the disks in the disk group a specific identifier.

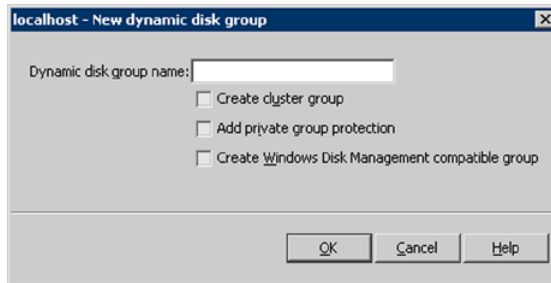
This identifier precedes a system-generated numeric ID that SFW uses to distinguish the disks in the disk group.

For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Click **Next** to continue.



If you want to create a new dynamic disk group to add the disk to, click **New dynamic disk group**. The New dynamic disk group window appears to let you indicate the disk group name and characteristics.



Details on the options for creating a disk group are available in the procedure on creating a dynamic disk group.

See [“Creating a dynamic disk group”](#) on page 44.

- 4 A confirmation screen appears listing the selected disk or disks. Click **Next** to continue if you are satisfied with the selections, or choose **Back** to go back to modify your disk selection.

For each disk, information is provided on whether it contains volumes and whether it is updated. If it is not going to be updated, you will want to modify the selection.

- 5 Click **Finish** in the final screen.

You can also create a new disk group for the new storage.

See [“About creating dynamic disk groups”](#) on page 43.

You can also remove a disk from a disk group.

See [“Remove a disk from a dynamic disk group”](#) on page 102.

Create new volumes

At this point, you are ready to create new volumes on the storage.

See [“About creating dynamic volumes”](#) on page 47.

Disk tasks

This section describes the commands that deal with disk changes and maintenance.

See [“Remove a disk from a dynamic disk group”](#) on page 102.

See [“Remove a disk from the computer”](#) on page 103.

See [“Offline a disk”](#) on page 104.

See [“Update disk information by using rescan”](#) on page 104.

- See [“Set disk usage”](#) on page 105.
- See [“Evacuate disk”](#) on page 105.
- See [“Replace disk”](#) on page 106.
- See [“Changing the internal name of a disk”](#) on page 107.
- See [“Renaming an enclosure”](#) on page 107.
- See [“Work with removable media”](#) on page 108.
- See [“View disk properties”](#) on page 110.
- See [“Add a new disk to your system”](#) on page 98.
- See [“Add a disk signature to a disk”](#) on page 98.
- See [“Disk monitoring”](#) on page 193.
- See [“Merge foreign disk command”](#) on page 538.

Remove a disk from a dynamic disk group

In the process of adding new storage, you may want to rearrange your disks in your existing disk groups. You can remove a dynamic disk from a dynamic disk group if there are no volumes on the disk. If volumes exist, you must delete them before you can remove the disk. Disks shown in Storage Foundation for Windows to be Unreadable or Foreign can be removed, but doing so causes you to lose all existing data on the disk that is removed from the group. If a disk has failed and is marked as Missing, you can also remove it from the dynamic disk group. However, you cannot remove the last disk from a disk group that is a Microsoft Failover Clustering resource.

To remove a disk from a dynamic disk group, use the Remove Disk from Dynamic disk group command. Once you remove all the disks from a dynamic disk group, the disk group is removed from the GUI, and the disks become basic disks.

Warning: Removing a disk that is Unreadable or Foreign from a dynamic disk group changes the disk into a basic disk and all of its data is lost.

To remove a disk from a dynamic disk group

- 1 Select the disk of the dynamic disk group and select **Remove Disk from Dynamic Disk Group**.
- 2 Make sure the disk or disks that you want to remove show in the right pane of the window.

Click **OK** to complete the operation.

Remove a disk from the computer

This section has information on removing a basic or dynamic disk from a computer.

Identifying a physical disk

To identify a physical disk that is represented by a disk that appears in the VEA GUI, use the Ping Disk command. This command flashes the fault light that is built into the housing of the physical disk. The command flashes the fault light until you stop the command. This is a very convenient feature when trying to identify a particular physical disk among many others.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Ping Disk command.

To identify a physical disk

- 1 Right-click the selected disk and select **Ping Disk** from the context menu.
- 2 In the screen that appears, click **Start** to begin pinging the physical disk. The fault light flashes until you click **Stop**.
- 3 Click **Close** when you are finished using the command.

Basic disk

If you remove a basic disk from the computer and then do a rescan, the disk and its volumes do not show up in the GUI.

Dynamic disk

If the status of a dynamic disk remains Offline and Missing and you determine that the disk has a problem that cannot be repaired, you can remove the disk from the system by using the Remove Disk from Dynamic disk group command.

However, before you can remove the disk, you must delete all volumes on the disk. You can save any mirrored volumes on the disk by removing the mirror that is on the Missing disk instead of the entire volume. Deleting a volume destroys the data in the volume, so you should remove a disk only if you are certain that the disk is permanently damaged and unusable.

If a disk is marked Missing and is not connected to the computer, you can use the Remove Disk from Dynamic disk group command without worrying about the data.

Once you have removed the disk from the disk group, the disk becomes a basic disk and it can then be removed from the computer. Complete the process by doing a rescan.

See [“Remove a disk from a dynamic disk group”](#) on page 102.

Note: If you have a disk that appears to be deteriorating because it has many errors, you can easily move the data from that disk to another disk by using SFW’s Move Subdisk command.

See [“Moving subdisks”](#) on page 244.

Offline a disk

On Windows Server operating systems, a basic disk can be taken offline for maintenance by using the Offline command.

Right-click a basic disk in the GUI to bring up a context menu and select the Offline command.

Use the Online command to bring the basic disk online.

Right-click the basic disk to bring up the context menu and select the Online command.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Offline command. However, the Online command can be used on dynamic disks belonging to a Microsoft Disk Management disk group.

Update disk information by using rescan

When Storage Foundation for Windows rescans disks, it scans all disks for disk configuration changes. It also updates information about CD-ROM drives and other removable media, basic volumes, file systems, and drive letters.

Rescanning disks can take several minutes, depending on the number of hardware devices installed.

Note the following:

- You should rescan disks when you remove disks from or add disks to a computer.
- From the Actions menu, select Rescan.
If you click the Tasks tab at the lower left corner of the VEA window, information on the progress of the rescan is displayed in the window’s lower pane.
More information is available.
See [“Additional important information about the rescan command”](#) on page 537.

Note: The Rescan command is only effective on the StorageAgent node and any of its sub-nodes.

Set disk usage

This command lets you make settings for the hot relocation function, including making a disk a hot spare.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support setting disk usage.

Note: Dynamic disks belonging to a SAN disk group do not support setting disk usage.

It is described in detail in the section about Hot Relocation.

See [“Hot relocation”](#) on page 216.

Evacuate disk

The Evacuate Disk command moves the entire contents of a healthy disk to the free space on another dynamic disk. Even if there is a failed volume on the original disk, the volume is moved successfully, but the data is not copied. You can restore the data backup on the moved volume, after it is reactivated.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Evacuate Disk command.

To evacuate a disk

- 1 Select the disk to evacuate.
Right-click the disk.
- 2 Select **Evacuate Disk**.

- 3 The Evacuate Disk dialog appears. Select either **Auto Assign destination disk** or **Manually assign destination disk**. If you chose **Manually assign destination disk**, assign the destination disk by selecting one disk from the display.

You may also check **Disable Track Alignment** to disable track alignment on the destination disk.
- 4 Click **OK** to evacuate the disk.

Replace disk

Disks most often need replacing when they fail or start to behave strangely. You can replace a failed dynamic disk by an empty basic disk. The volume configuration is recreated on the new disk. Active redundant volumes (mirrored and RAID-5) on the replaced disk are resynchronized automatically after disk replacement.

Failed volumes on the replaced disk, however, stay failed after disk replacement because no valid data is available for copy.

The disk replacement procedure can be performed only on a disk that has failed. The VEA console identifies the disk by renaming it "Missing Disk." If the disk replacement is successful, the replacement disk takes on the attributes of the failed disk, including the disk name.

Note: This option is available only on Missing Disks.

If the original disk is returned to the system after it was replaced, it shows in the VEA console as a foreign disk and a disk group called Unknown Dg is created. Use the Merge Foreign Disk command to return the disk to the same dynamic disk group it was in before.

See [“Merge foreign disk command”](#) on page 538.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Replace Disk command.

To replace a disk

- 1 Right-click the Missing Disk.
- 2 Select **Replace Disk**.

A dialog box appears containing a list of empty basic disks.
- 3 Select the disk you want to use to replace the Missing Disk.
- 4 Click **OK** to perform the disk replacement.

Changing the internal name of a disk

You can change the internal name of a disk with the VEA GUI to make it easier to identify.

The internal name of a disk can also be changed with the `vxedit` CLI command.

See “[vxedit](#)” on page 687.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support changing the internal name of a disk.

To change the internal name of a disk

- 1 Expand the Disks node in the tree view.
- 2 Right-click the desired disk node.
- 3 Select **Change Disk Internal Name** from the context menu.

A dialog box appears to rename the disk.

- 4 Enter the new internal name of the disk.
- 5 Click **OK** to complete the operation.

Renaming an enclosure

You can rename an enclosure with the VEA GUI to make it easier to identify.

An enclosure can also be renamed with the `vxedit` CLI command.

See “[vxedit](#)” on page 687.

SFW provides an enclosure node in the tree view to represent the enclosure for the internal disks of your server. This enclosure cannot be renamed.

Renaming an enclosure in a cluster environment only renames the enclosure on the node that the rename operation occurs. You need to rename the enclosure on each of the other nodes in the cluster to maintain the same enclosure name throughout the cluster.

To rename an enclosure

- 1 Expand the Enclosures node in the tree view.
- 2 Right-click the desired enclosure node.
- 3 Select **Rename Enclosure** from the context menu.

A dialog box appears to rename the enclosure.

- 4 Enter the new name of the enclosure.
- 5 Click **OK** to complete the operation.

Work with removable media

Removable media include CD-ROM disks and other removable disks, such as Zip disks. There are a limited number of actions you can do in Storage Foundation for Windows for removable disks.

You can perform the following actions:

- You can create only one primary partition on removable-media devices. You cannot create extended partitions, logical drives, or dynamic volumes on removable-media devices. The primary partition must use up all unallocated space on the removable disk.
- You cannot upgrade removable-media devices to dynamic volumes.
- Removable-media devices can be formatted as FAT, FAT32, or NTFS.
- You cannot mark the primary partition on a removable-media device as active.
- You cannot delete the primary partition on a removable-media device.
- The Eject command lets you eject a removable-media disk through the software rather than by a manual ejection mechanism.

Working with disks that support thin provisioning

Thin provisioning is a technology to allocate disk storage on an as needed basis. Disks that support thin provisioning are disks that can allocate storage as data is written to the disk. In addition, some of these disks can reclaim unused storage on the disk.

Disks that support thin provisioning have the following characteristics:

- Disk storage allocated as data is written to the disk.
- Unused storage may be reclaimed to minimize the amount of unused space on the disk.

Note: Not all disk vendors support storage reclamation. Disks that support both thin provisioning and storage reclamation are represented in the VEA GUI as a disk with a green colored sector with an asterisk (*). Disks that support only thin provisioning are represented in the VEA GUI as a disk with a red colored sector.

Thin provisioning and reclamation is enabled on TRIM and UNMAP-supported devices. UNMAP is the default reclamation interface for devices that support thin provisioning, except for SSDs. For SSDs that support TRIM, the default reclamation interface is TRIM.

Reclaiming storage space from thin provisioned disks can be done for disk groups, disks, or volumes. The table briefly describes SFW wizards and CLI commands that support reclaiming unused storage.

Table 4-1

Wizard	CLI Command	Description of Operation
Disk group Reclaim Wizard	<code>vx dg reclaim</code>	This wizard and CLI command can reclaim unused storage from the volumes and disks that reside in the disk group. Each can be set up to reclaim storage space from volumes; to reclaim free space from disks; or to reclaim from both volumes and disks that reside in the disk group.
Disk Reclaim Wizard	<code>vx disk reclaim</code>	This wizard and CLI command can reclaim unused storage from a disk.
Volume Reclaim Wizard	<code>vx vol reclaim</code>	This wizard and CLI command can reclaim unused storage from a volume.

The timeout value for the reclaim I/O operations is set to 30 seconds by default. However depending on the configuration, the storage reclaim operation may sometimes take a significant amount of time to complete.

To avoid the reclaim I/O failures due to timeout errors, you can change the reclaim I/O timeout value by adjusting the following registry tunable parameter:

```

HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager
\MaxReclaimIOTimeout\VolumeManager\MaxReclaimIOTimeout

```

Thin provisioning reclaim operation is blocked on the following:

- RAID-5, DCO, or Volume Replicator DCM volumes
- FAT, FAT32, or raw dynamic volumes

- Volumes configured for snapshots

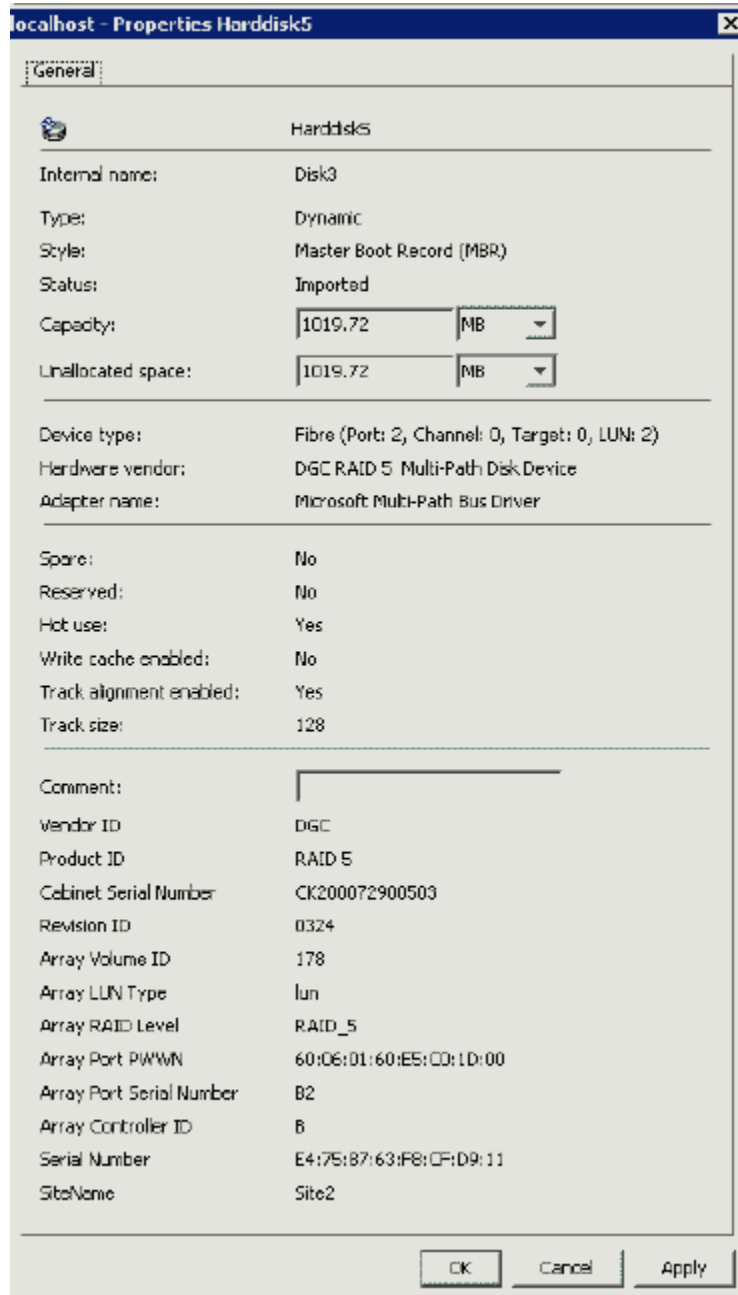
Note: The reclaim storage operation is disabled during a shrink volume operation. The reclaim storage operation is enabled when the shrink volume operation has completed.

View disk properties

Right-click a disk and select Properties from the context menu. The categories for the properties of a disk differ, depending on whether the disk is basic or dynamic and whether it is directed or part of an array.

The screen that follows shows an example of a Properties window for a dynamic disk.

Figure 4-1 Disk Properties



The categories in the sample screen are described in the table that follows.

Table 4-2 Disk Properties Categories

Category	Description
Internal Name	The name of the hard disk is shown at the top of the window.
Type	Fields are Basic and Dynamic.
Style	Refers to a partition style. The two categories are Master Boot Record (MBR) and GUID Partition Table (GPT).
Status	Fields are Imported and Deported.
Capacity	Size in Sectors, KB, MB, GB, TB.
Unallocated space	Amount of free space that is available to create a partition or volume.
Thin Type	Indicates the type of thin provisioned disk, Thin, or Thin Reclaim. Note: For thin provisioned disks only.
Provisioned size	Current amount of storage allocated to the disk that supports thin provisioning. Note: For Thin Reclaim provisioned disks only. Not all disks that are enabled for thin provisioning provide the provisioned size value.
Allocation unit size	Minimum amount of storage that is added to a thin provisioned disk to meet current storage requirements. Note: For Thin Reclaim provisioned disks only.
Device type	The Device type category includes SCSI, Fibre, and some other less commonly used types. If a controller card is involved, gives port, channel, target, and LUN information.
Hardware vendor	Vendor information from the disk itself.
Adapter name	Gives information on the host bus adapter.
Spare	Indicates whether the disk is a preferred hot relocation target.
Reserved	Indicates whether the disk is reserved—that is, excluded from use in automatic selection operations.
Hot use	Indicates whether the disk can be used to receive volumes from the hot relocation procedure.
Write cache enabled	Indicates whether the write cache is enabled.

Table 4-2 Disk Properties Categories (*continued*)

Category	Description
SCSI-3 Reserved	Indicates whether the disk is enabled for SCSI-3 reservations.
Track alignment enabled	Indicates whether the disk is enabled for track alignment.
Track size	Indicates the size of a track on the disk.
Comment	User-defined comment.
SiteName	If disk is site tagged to a host, then the name of the site is displayed to which it is added.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support disk comments.

You can also view disk characteristics through the Disk View in the GUI and through the GUI's right-pane view.

See [“The disk view”](#) on page 77.

Veritas Disk ID (VDID)

The Veritas Disk ID (VDID) is a string that contains information that can identify the array that a disk is contained in.

The VDID string is shown in the GUI as a part of the display of information for disks. The VDID string is also shown as a result of the `vxdisk diskinfo` and the `vxdisk list` CLI commands.

The format of the VDID string varies by array manufacturer and array model. The details of how to interpret the VDID string for a particular array are in an appendix.

See [“Format of VDID strings for disks belonging to various arrays”](#) on page 757.

Generic VDID

Veritas has a generic VDID for the arrays that are not supported by the VDID library. For such arrays, the generic VDID string helps identify and categorize disks belonging to various arrays using the Unique Disk ID. In the VEA GUI, all the disks for these arrays are gathered under a common **Enclosure** named as `Disks@<HostName>`.

Similar to the VDID string, information about the generic VDID string can be viewed by using the VEA GUI (displayed as part of information for disks) and the `vxdisk`

`diskinfo` and `vxdisk list` CLI commands. The format of the generic VDID string for disks belonging to various arrays is as follows:

```
<VendorID>_<ProductID>_DISKS_<LUNSerialNumber>
```

General Partition/Volume tasks

This section describes the general partition and volume tasks.

See [“Delete a volume”](#) on page 115.

See [“Delete a partition or logical drive”](#) on page 115.

See [“Shredding a volume”](#) on page 115.

See [“Refresh drive letter, file system, and partition or volume information”](#) on page 116.

See [“Renaming a mirror \(plex\)”](#) on page 118.

See [“Changing the internal name of a volume ”](#) on page 119.

See [“Mount a volume at an empty folder \(Drive path\)”](#) on page 119.

See [“View all drive paths”](#) on page 120.

See [“Format a partition or volume with the file system command”](#) on page 121.

See [“Cancel format”](#) on page 121.

See [“Change file system options on a partition or volume”](#) on page 121.

See [“Set a volume to read only”](#) on page 122.

See [“Check partition or volume properties”](#) on page 122.

See [“Expand a dynamic volume”](#) on page 124.

See [“Expand a partition”](#) on page 128.

See [“About creating dynamic volumes”](#) on page 47.

See [“Create primary and extended partitions ”](#) on page 135.

See [“New logical drives”](#) on page 137.

See [“Subdisk move, split, and join”](#) on page 241.

See [“Set up a dynamic boot and system volume \(Optional\)”](#) on page 59.

See [“Bring a basic volume back to a healthy state”](#) on page 532.

See [“Bring a dynamic volume back to a healthy state”](#) on page 533.

Delete a volume

Perform the following steps to delete a volume.

Warning: All data in the volume is lost when you delete it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

To delete a volume

- 1 Select the designated volume.
Select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete or **No** to cancel. The volume is removed immediately if you click **Yes**.

Delete a partition or logical drive

Perform the following steps to delete a partition or logical drive.

Warning: All data in the partition or logical drive is lost when you delete it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

To delete a primary partition, an extended partition, or a logical drive

- 1 Select the partition or logical drive you want to delete and choose the command to delete the partition or logical drive from the context menu.
- 2 You are prompted for verification. Click **Yes** to delete or **No** to cancel. The partition or logical drive is removed immediately if you click **Yes**.

Note that you cannot delete a partition with the system files (the boot volume).

Shredding a volume

Shredding a volume destroys the data stored on a volume by overwriting the volume with a digital pattern.

SFW provides the following methods for overwriting the volume:

- 1 Pass algorithm Overwrites the volume with a pattern of zeroes.

- 3 Pass algorithm** The volume is overwritten a total of 3 times. The volume is first overwritten with a pre-selected digital pattern, then overwritten with the binary complement of the pattern. In the last pass, the volume is overwritten with a randomly selected digital pattern. This 3 pass approach is based on the specifications in the US Department of Defense standard 5220.22-M (3).
- 7 Pass algorithm** The volume is overwritten a total of 7 times. Each pass consists of overwriting the volume with a randomly selected digital pattern or with the binary complement of the previous pattern. This 7 pass approach is based on the specifications in the US Department of Defense standard 5200.28-STD (7).

Warning: Operations performed on logs are done on the volumes that they are associated with. All the data in the volume is 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.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Shred Volume command.

To shred a volume

- 1** Right-click the designated volume in the VEA GUI and select **Shred Volume**.
- 2** Select the method to shred the volume and click **OK** to continue.
- 3** In the confirmation dialog box, click **Yes** to confirm the operation.

The volume is entirely overwritten and removed immediately when the operation has completed.

Refresh drive letter, file system, and partition or volume information

To refresh, perform the following steps:

To refresh drive letter, file system, and partition or volume

- 1 To refresh system information, select **Refresh** from the **Actions** menu

The command refreshes drive letter, file system, volume, and removable-media information on the current computer. It also checks to see whether previously unreadable volumes are now readable. It may not pick up the disk changes that occurred since the last restart or rescan if there was no I/O occurring on the changed disk.

- 2 Click the **Tasks** tab in the lower pane to see a progress bar that displays the percentage of completion for the refresh process.

Note: The Refresh command is only effective on the StorageAgent node and any of its sub-nodes.

Add, change, or remove a drive letter or path

You cannot change the drive letter or path of the active volume (the computer's boot volume). In previous versions of Volume Manager for Windows, the command Change Drive Letter and Path was on the context menu for a volume or partition. Now it is a subcommand under the File System command that is on the volume or partition context menu.

Consider the following:

- If you try to assign a drive letter to a volume that is in use, a dialog box asks you whether you want to forcibly reassign the drive letter. You can click **Yes to do so**.
- You cannot modify an existing drive path in this window. You have to remove the drive path and recreate it. Select the Remove option. Then you have to invoke the command again to recreate the drive path with the Add button.

Warning: Changing the drive letter or path of a partition or volume may cause programs installed on that partition or volume to no longer run.

To add, change, or remove a drive letter or path

- 1 Right-click the designated partition or volume and on the context menu, select **File System > Change Drive Letter and Path**.

The Drive Letter and Paths screen appears.

- 2 Select **Add**, **Modify**, or **Remove**, depending on what you want to do.

Choose from the following:

- If a partition or volume does not have a drive letter or path, select **Add**.
 - If you want to change an existing drive letter or path, select **Modify**.
 - If you want to remove a drive letter or drive path, select **Remove**.
- 3** Select **Add** to add a drive letter or path to a partition or volume.
- You can select **Assign a drive letter** or **Mount as an empty NTFS folder**.
- Choose from the following:
- The **Assign a drive letter** option lets you select a drive letter from the drop-down list. Click **OK** to assign the drive letter after you have made your selection.
 - The **Mount as an empty NTFS folder** option lets you specify the folder. To specify the folder, click the **Browse** button. The Browse for drive path window appears: Navigate to an existing folder or create a new folder by using **New folder**. Click **OK** to mount the volume. See “[Mount a volume at an empty folder \(Drive path\)](#)” on page 119.
- 4** Select **Modify** to modify the drive letter. The **Assign a drive letter** box lets you specify the drive letter. Select a new drive letter from the drop-down list, and click **OK** to assign the drive letter
- 5** Select **Remove** to remove a drive letter. Click **OK** to continue.
- Click **Yes** in the confirmation screen that appears to complete the operation.

Renaming a mirror (plex)

You can change the name of a mirror to make it easier to identify.

The name of a mirror (plex) can also be changed with the vxedit CLI command.

See “[vxedit](#)” on page 687.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support the Rename Plex command.

To change the name of a mirror

- 1** Expand the Volumes node in the tree view.
- 2** Select the desired volume that contains the mirror that you want to rename.
- 3** Click the **Mirrors** tab in the right pane of the VEA GUI.

- 4 Select the desired mirror and select **Rename Plex** from the context menu.
A dialog box appears to rename the mirror (plex).
- 5 Enter the new name of the mirror (plex).
- 6 Click **OK** to complete the operation.

Changing the internal name of a volume

You can change the internal name of a volume to make it easier to identify.

The internal name of a volume can also be changed with the `vxedit` CLI command.

See “[vxedit](#)” on page 687.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support changing the internal name of a volume.

To change the internal name of a volume

- 1 Expand the Volumes node in the tree view.
- 2 Right-click the desired volume node.
- 3 Select **Change Volume Internal Name** from the context menu.
A dialog box appears to rename the volume.
- 4 Enter the new internal name of the volume.
- 5 Click **OK** to complete the operation.

Mount a volume at an empty folder (Drive path)

Storage Foundation for Windows lets you mount a volume at any empty folder on a local NTFS volume. SFW assigns a drive path to the volume rather than a drive letter. Drive paths are useful because they eliminate the 24-drive-letter limit on hard-disk volume names. (Drive letters A and B are reserved for floppy drives.) The volume can be a partition, a logical drive that was created in Disk Administrator, or a dynamic volume.

For example, you can mount the C:\Temp folder as another drive to provide additional disk space for temporary files.

The only restrictions are that the folder you mount the volume to must be empty and it must be on an NTFS volume on a local computer. You cannot mount volumes to network-accessible drives. You can, however, assign multiple drive paths to a single volume.

Mounted drive folders are distinguished from ordinary folders by having a drive icon instead of the normal folder icon.

You can assign the drive path when you create the volume with the New Volume wizard at the Assign Drive Letter screen. You can also use the command Change Drive Letter and Path after the volume is created. A volume that already has a drive letter assigned can also have one or more drive paths. It does not have to be an unlabeled volume.

The Change Drive Letter and Path command is used to mount a volume.

See [“Add, change, or remove a drive letter or path”](#) on page 117.

See [“View all drive paths”](#) on page 120.

Unmount a volume

To unmount a volume with a drive path, use the command Change Drive Letter and Path, select the drive path, and click Remove.

Note: To modify a drive path, remove it and then create a new drive path using the new location. You cannot modify the drive path directly.

View all drive paths

This command lets you view all volumes that have a drive path on a selected server.

To view all mounted volumes

- 1 Right-click the computer icon and select **View All Drive Paths**. In addition, if the computer icon is already selected, you can choose **View All Drive Paths** from the **Actions** menu.

The View All Drive Paths window appears.

- 2 If you want to unmount any volume listed in this window, select the drive path and click **Remove Drive Path**.

You can also use the **Change Drive Letter and Path** command to remove a drive path.

- 3 Click **OK**.

See [“Mount a volume at an empty folder \(Drive path\)”](#) on page 119.

Format a partition or volume with the file system command

If you want to format a volume after it has been created, you need to select the File System command from the context menu and then choose the Replace File System subcommand.

Warning: All data in the partition or volume is lost when you format it.

To format a partition or volume

- 1 Right-click the volume or partition you want to format, click **File System** from the context menu, and then click the subcommand **Replace File System**.
- 2 Choose NTFS, ReFS, FAT, or FAT32 as the file system to use.
- 3 Select an allocation size, or use the default which is automatically selected.
- 4 Enter a label for the volume.

This label appears on the VEA console. If a name has been selected, this name appears in this text box, but you can change it here.

- 5 (Optional) Select the following formatting options:
 - Perform a quick format — Formats the volume or partition without scanning for bad sectors in the volume or partition.
 - Enable file and folder compression — Can be used only if you choose NTFS format.

Click **OK**.

A confirmation screen appears, warning that formatting erases data.

- 6 Click **OK** in the confirmation screen to begin formatting.

You can view the percentage of completion for the formatting process in the bottom pane of the console on the **Tasks** tab.

Cancel format

The Cancel Format subcommand is a third option from the File System menu. This command appears only during the time the file system is being formatted and gives the user the option to cancel the formatting operation.

Change file system options on a partition or volume

When a partition or volume is created, a file system options screen is included with the Create Partition or New Volume wizard. Once the partition or volume is created,

you can access the same file system options screen by using the File System command and then make desired changes to the options.

To change file system options

- 1 Right-click the partition or volume and choose the commands **File System>Replace File System** from the context menu.
- 2 Once the File System screen appears, change the desired attributes.
You may want to format the partition or volume.

See [“Format a partition or volume with the file system command”](#) on page 121.

Set a volume to read only

Perform the following steps to set a volume to read only.

Note: Changing a volume from read/write to read only may affect the applications that access the volume.

To set a volume to read only

- 1 Right-click the volume to be set to read only.
- 2 Choose **Properties** from the context menu that appears.
The Properties window comes up.
- 3 Check the **Read only** check box.
- 4 Click **OK**.

Check partition or volume properties

You can check the properties of a partition or volume by right-clicking a partition or volume and selecting Properties from the context menu. The Properties window comes up. The categories for the properties of a volume differ, depending on the volume type.

The screen that follows shows the Properties window for a Mirrored Concatenated volume.

Figure 4-2 Volume Properties



The categories for the Mirrored Concatenated volume shown in the sample screen are as described in the following table:

Table 4-3 Volume Properties Categories

Category	Description
Name	The volume name is shown at the top of the window.
Layout	Indicates the volume layout type. The layout categories are concatenated, striped, RAID-5, mirrored, and mirrored striped.
Status	Status of the volume.
Copies	Number of mirrors the volume has.
Columns	Number of columns. It is two or more for striped volumes.
Log	Whether there is a DRL or RAID-5 log or FastResync
Group name	Basic or dynamic disk group the volume is associated with.

Table 4-3 Volume Properties Categories (*continued*)

Category	Description
Device path	The path to the device.
Size	Size of the volume in Sectors, KB, MB, GB, TB.
Comment	User-defined comment.
File system type	Fields are FAT, FAT 32, NTFS, ReFS.
File system label	Operating system file system label.
File system capacity	Size of file system.
File system free space	Amount of free space available in file system.
Aligned	Whether the volume is track aligned or not.
Sitetype	Indicates whether the volume is Siteless, Site Confined, or Site Separated.
Attributes	Check box to enable volume as read only.

Note: Dynamic disks belonging to a Microsoft Disk Management disk group do not support volume comments.

Expand a dynamic volume

The Expand Volume command lets you increase a dynamic volume's size. You can have the program automatically assign destination disks, or you can manually assign these disks.

You can expand a volume only in case of the following:

- The volume is formatted with NTFS, ReFS, or is a raw dynamic volume
- There is unallocated space on a dynamic disk within the dynamic disk group onto which the volume can be extended.

You cannot expand a volume if one of the following is true:

- The volume is formatted with FAT or FAT32.
- The volume is set to read only.
- There is not enough unallocated space available on the dynamic disks within the dynamic disk group to extend the volume.

You can extend volumes onto a maximum of 256 disks. No portion of an extended volume can be deleted without deleting the entire volume.

A system or boot volume is extended in increments of the disk's cylinder size and only into contiguous space at the end of the volume.

More information is available if a dynamic boot or system volume fails to extend.

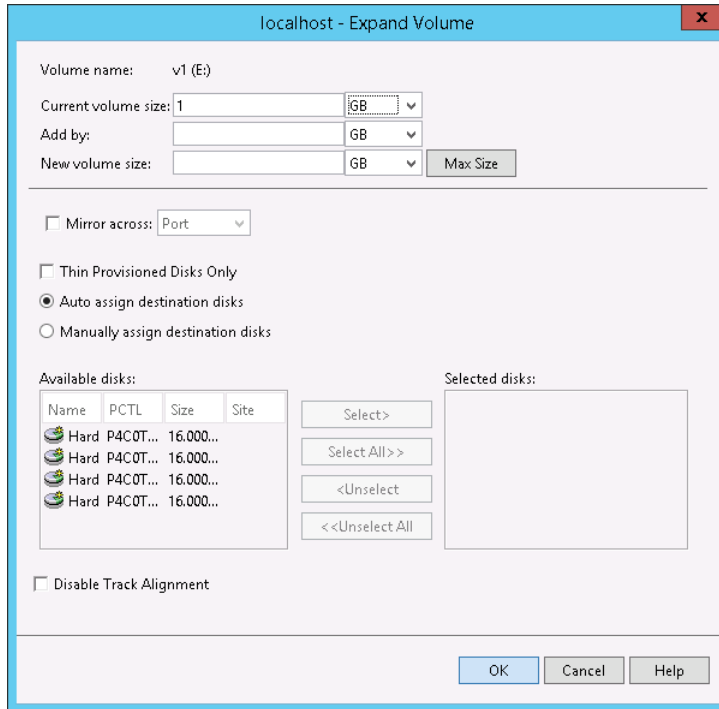
See [“Cannot extend a dynamic boot or system volume”](#) on page 548.

For Site-based volumes, auto-grow and manual grow operations always adhere to site boundary restrictions.

See [“Administering volumes for site-based allocation”](#) on page 456.

To expand a dynamic volume

- 1 Right-click the dynamic volume you want to expand and select **Resize Volume** > **Expand volume** from the context menu that appears.



2 Complete the following on the **Expand Volume** panel:

Current Volume Size	Specifies the current size of the volume. Select from the drop-down list volume size in Sectors, KB, MB, GB, or TB.
Add by	This option lets you add a specific value to current volume by entering a numerical value in the box. Select the drop-down list to indicate volume size in Sectors, KB, MB, GB, or TB.
New volume size	Specifies the total volume size after growing the volume.
Max size	Click Max Size to enter the maximum volume size. Select the drop-down list to indicate volume size in Sectors, KB, MB, GB, or TB.
Auto assign destination disks	Select this option to assign disk or disks automatically.
Mirror across	Select this check box to mirror across disks by Port, Target, Enclosure, or Channel.
Thin Provisioned Disks Only	Select this check box to list only thin provisioned disks under the "Available disks" box. If not selected, then all disks are listed.
Manually assign destination disks	Select this option to assign disk or disks manually. Use the Select or Select All buttons to select disk or disks respectively
Disable Track Alignment	Select this check box to disable track alignment for the volume.

Click **OK**.

Expanding a volume that has a snapshot

If you expand a volume that has a snapshot or detached mirror, then you also need to consider increasing the size of the DCO volume that tracks the changes made

to the volume. Increasing the size of the DCO volume allows for a more efficient mirror resynchronization process.

Before expanding the volume, you should perform snap back and snap abort operations on the volume. The snap abort operation should be done with the Convert into normal mirror option. This deletes the existing DCO volume and the snap volume becomes the mirror volume with a DCO volume. Right-click the volume and Select remove log to remove the DCO volume. After you expand the volume and prepare the volume for a snapshot with the Select existing mirror for snap option, a new DCO volume is created automatically at the optimal size for the expanded volume.

More information is available about snapshots.

See [“Snapshot commands”](#) on page 259.

More information is available about DCO volumes.

See [“Disk Change Object \(DCO\) volume”](#) on page 256.

Expand a partition

With SFW, the capability has been added to expand a partition (also known as a basic volume). The command is done online and is similar in concept to the command to expand a dynamic volume.

To expand a partition (basic volume)

- 1 Right-click the basic volume you want to expand, and then click **Expand Partition** in the context menu that appears.
- 2 Enter a number in the desired text box.

The options are as follows:

- The **Add by** option lets you add a specific amount to the partition by entering a number in the entry box.
- The **New Partition size** option lets you specify the total partition size you want in the entry box.

Clicking the **Max Size** button enters the maximum partition size in the **New Partition size** entry box.

The drop-down lists next to the entry boxes let you indicate the size in Sectors, KB, MB, GB, or TB.

- 3 Click **OK** to complete the command.

Safeguarding the expand volume operation in SFW against limitations of NTFS

To prevent users from expanding volumes beyond the limitations of NTFS for NTFS cluster size, the following safeguard is implemented:

- SFW does a check for the NTFS cluster size of the volume while expanding it.
- The operation fails, if the user tries to expand the volume beyond a limit.
- The limit depends upon the cluster size of the specified volume during its creation.

When SFW receives an **expand volume** request, it validates the requested new size against the **MAX SIZE** supported for that volume.

If the size exceeds the limit, the **expand operation** results in a failure, displaying the following error message:

The extend operation did not complete because the cluster count is higher than expected.

Architecturally NTFS supports: 2^{32} clusters minus one cluster count.

The General formula is:

$((2^{32})-1) * (\text{user-selected size in bytes of a single cluster before format})$

= maximum FS size for that instance possible during grow/extend without reformat.

So we have $((2^{32}) - 1 = (4,294,967,296) - 1 = 4,294,967,295$ clusters count.

Therefore, for a cluster size of 512 bytes, maximum size of the mountable file system:

$(2^{32})-1$ cluster count X 512 bytes represented by each cluster = 2199023255040 bytes.

Hence, for a 512 Bytes cluster sized volume, (2199023255040 bytes) 2 TB is the supported file system size.

The following table is used to check limitations against the cluster size:

Table 4-4 Limitation against the cluster size

Cluster Size	Maximum NTFS Volume Size
512	2,199,023,255,040 (2TB)
1,024	4,398,046,510,080 (4TB)
2,048	8,796,093,020,160 (8TB)
4,096	17,592,186,040,320 (16TB)

Table 4-4 Limitation against the cluster size (*continued*)

Cluster Size	Maximum NTFS Volume Size
8,192	35,184,372,080,640 (32TB)
16,384	70,368,744,161,280 (64TB)
32,768	140,737,488,322,560 (128TB)
65,536	281,474,976,645,120 (256TB)

Safeguarding the expand volume operation in SFW against limitations of ReFS

To prevent users from expanding volumes beyond the limitations of ReFS for ReFS cluster size, the following safeguard is implemented:

- SFW does a check for the ReFS cluster size of the volume while expanding it.
- The operation fails, if the user tries to expand the volume beyond a limit.
- The limit depends upon the cluster size of the specified volume during its creation.

When SFW receives an **expand volume** request, it validates the requested new size against the **MAX SIZE** supported for that volume.

If the size exceeds the limit, the **expand operation** results in a failure, displaying the following error message:

The extend operation did not complete because the cluster count is higher than expected.

Architecturally, ReFS supports a single cluster size of 64KB.

Shrink a dynamic volume

You can decrease or shrink the size of a dynamic volume using the online volume shrink feature. The Shrink Volume command is helpful in reclaiming unused space to better use your resource. The command supports all dynamic volume types, such as simple, spanned, striped, mirrored, and RAID-5 volumes.

Note: The shrink volume operation is disabled during a reclaim storage operation for thin provisioned disks. The shrink volume operation is enabled when the reclaim storage operation has completed.

About the Shrink Volume command

The Shrink Volume command lets you decrease the size of a dynamic volume. The command calculates the amount of space that can be freed from the volume to create a new smaller volume size. The size of a volume after the shrink volume operation is approximately the difference of the current volume size and the number of maximum reclaimable bytes. The new volume size is displayed in the Veritas Enterprise Administrator (VEA) GUI.

During the shrink volume operation, data may be relocated on the volume so that the volume can be made into the smaller specified size.

Before shrinking a dynamic volume

Consider the following before shrinking a dynamic volume.

- Before performing the volume shrink operation, you must install the KB 2615327 hotfix from Microsoft.
- If the combined length of the volume name and disk group name is more than 9 characters, then you must install the KB 2619083 hotfix from Microsoft before shrinking the volume.
- You can shrink a volume only if the dynamic volume is either on raw partition (that is, without a file system) or using the NTFS file system.
- The volume must be online.
- If a volume or degraded volume has any snapshots or detached mirrors, then the volume shrink operation is rejected. To shrink a volume that has been snapshotted, you must first snapback the volume and then shrink the volume. After shrinking a volume, you can snapshot the volume again.
- You can neither shrink a volume of less than 1 MB in size nor shrink a volume by less than 1 MB.
- You cannot shrink a volume beyond the point where any unmovable files, such as the page file or the shadow copy storage area, are located. Unmovable files are not automatically relocated when you shrink a volume.
If you need to shrink the volume further, move the page file to another disk and delete the stored shadow copies, shrink the volume, and then move the page file back to the disk.
- If dynamic bad-cluster remapping detects several bad clusters, then you cannot shrink the volume. If this occurs, then you should move the data and replace the disk.
- Shrink Volume does not support read-only volumes.

- Volumes of more than 10 GB in size can be shrunk up to 3 GB approximately. In other words, the new size of the volume must be at least 3 GB approximately.
- Shrink Volume is not allowed during a mirror resynchronization, sub disk move, or thin reclamation.
- If a volume shrink fails after a successful file system shrink, it leaves the file system in shrunk state. However, the volume remains of the original size. For a workaround, refer to Known Issue 2385051 in the *Veritas InfoScale Release Notes*.
- The time taken by the volume shrink operation depends on the size of the volume and the time required to relocate the files.
- During the volume shrink operation, if the deport disk group operation is performed, either by the user or because of a failover, then the volume shrink operation is canceled.

Shrinking a dynamic volume

Follow the steps below to shrink a dynamic volume.

To shrink a dynamic volume

- 1 Right-click the dynamic volume that you want to shrink, click **Resize Volume**, and then select **Shrink**.
- 2 The **Shrink Volume** dialog box appears.
- 3 In the **Shrink Volume** dialog box, provide the volume size in the **Shrink by** box or **New volume size** box, or click the **Max Shrink** button.

The options offer the following:

- The **Shrink by** box lets you specify an amount by which you want to decrease the volume.
- The **New volume size** box lets you specify the new size of the volume.
- The **Max Shrink** button lets you decrease the volume by the maximum amount possible.

Note: Using the **Max Shrink** button with RAID-5 volumes may not yield optimal results due to the configuration of disks supporting the volume.

The drop-down lists next to the entry box let you indicate the amount in Sectors, KB, MB, GB, or TB.

After entering a value in the **Shrink by** box or clicking the **Max Shrink** button, the volume size that results from the Shrink Volume command is calculated and displayed in the **New volume size** box.

- 4 Click **OK** to start the volume shrink operation.

Canceling the shrink volume operation

Follow the steps below to cancel the shrink volume operation.

To cancel the shrink volume operation

- 1 Right-click the dynamic volume you want to cancel the shrink volume operation, highlight **Resize Volume**, and then click **Cancel Shrink** in the context menu that appears.
- 2 Click **OK** to confirm the cancelation and complete the command.

Dynamic LUN expansion

Many modern disk arrays allow existing LUNs to be resized. Veritas Volume Manager (VxVM) supports dynamic LUN expansion, by providing the Rescan utility.

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 should be performed on the host where the disk group is imported (or on the master node for a cluster-shared disk group).

When a disk is resized from the array side, you must perform the Rescan operation from the VxVM side to make VxVM aware of the new size. After the Rescan operation the volumes on the device should be grown.

Before shrinking a LUN, first shrink any volumes on the LUN or move those volumes off of the LUN. Finally, resize the LUN itself using the storage array's management utilities.

Basic disk and volume tasks

This section describes the procedures for working with basic disks and volumes that are under the control of SFW.

See [“What can you do with a basic disk?”](#) on page 134.

See [“Converting the partition style of a disk”](#) on page 134.

See [“Create primary and extended partitions ”](#) on page 135.

See [“New logical drives”](#) on page 137.

See [“Mark a partition as active”](#) on page 138.

See [“Change a basic disk to dynamic”](#) on page 138.

See [“Troubleshoot basic disks and volumes”](#) on page 139.

See [“Expand a partition”](#) on page 128.

What can you do with a basic disk?

Disks configured on a fresh system are defined as basic disks. You can then upgrade them to dynamic disks.

In general, a basic disk is managed with the Master Boot Record (MBR) or GUID Partition Table (GPT) partitioning scheme and may have partitions defined on it, but this is not required. The MBR style is supported in MS-DOS, Windows 95/98, and later Windows versions. MBR disks can contain up to four primary partitions or three primary partitions plus an extended partition. The extended partition may be further divided into logical drives. The GPT style allows a maximum of 128 primary partitions.

With the MBR style, you can also create new logical drives that reside in extended partitions. The logical drives are simple volumes that are limited to the space on the extended partitions. They cannot span multiple drives.

Basic disks can also contain RAID volumes that were originally created in Windows NT Disk Administrator, including simple and spanned volumes (volume sets), mirrored volumes (mirror sets), striped volumes (stripe sets), and RAID-5 volumes (stripe sets with parity). These volumes are also called "FT volumes" in Windows NT documentation.

SFW versions 5.0 and 5.1 do not support FT volumes. Only earlier versions of SFW provide support for FT volumes.

Versions of SFW earlier than SFW 5.0 let you maintain and repair these volumes but not to create new ones. Although you cannot create new NT Disk Administrator-type volumes on a basic disk, you can upgrade the existing volumes of this type to dynamic. Then these volumes have the full capabilities of other SFW dynamic volumes.

Converting the partition style of a disk

You can convert the partition style of a disk if the disk does not contain any data.

To convert the partition style of a disk

- 1 Right-click the basic disk to display the context menu. Select **Convert to MBR Disk**, if you have a GPT disk; or select **Convert to GPT Disk**, if you have an MBR disk.

The convert pop-up window appears.

- 2 Read the information in the pop-up window and click **Yes** to confirm the conversion.

Create primary and extended partitions

A partition is a portion of a physical disk that functions as if it were a physically separate disk. Partitions can be created only on basic disks. Partitions cannot span disks; they must be contiguous regions. An MBR disk can contain up to four primary partitions or three primary partitions plus an extended partition. The extended partition may be further divided into logical drives. A GPT disk can contain a maximum of 128 primary partitions.

Use the Create Partition wizard to create primary partitions and extended partitions on a basic disk.

To create primary and extended partitions

- 1 Right-click a basic disk to display the context menu, and select **New Partition**.

The Create Partition Wizard window appears.

- 2 Read the information on the first screen and then click **Next** to continue.

The Select disk and region window comes up.

- 3 Select a disk and free space on which to create the partition.

The dialog box shows a list of available disks and free spaces. Check the box for the disk and free space you want to use, and click **Next** to continue.

If there is no free space on the selected disk, the free space selection and the **Next** button are grayed out. Click **Cancel** to exit the wizard.

If the free space selection is grayed out, you cannot continue. This can occur if the disk already has the maximum number of partitions (either four primary partitions or three primary partitions and one extended partition). It also can happen if the available free space is in an extended partition, because only logical drives (not partitions) can be created there.

The Select Partition Type window comes up next.

- 4 Select a partition type (primary or extended) and indicate the size of the partition. The pull-down list to the right of the entry box for the disk space lets you select Sectors, KB, MB, GB, or TB. Click **Next** to continue.

The Assign Drive Letter screen comes up next.

- 5 Assign a drive letter or drive path to this volume if you want to format the partition. You can also assign a drive letter after creating the partition. Click **Next** to continue.

When you mount a local drive at an empty folder on an NTFS volume, the operating system assigns a drive path to the volume rather than a drive letter.

To mount a volume at an empty NTFS folder, click the radio button **Mount as an empty NTFS folder** and type the path to the empty folder, or click the **Browse** button to locate it.

See [“Mount a volume at an empty folder \(Drive path\)”](#) on page 119.

When creating an extended partition, no additional information is required and the final screen appears. Click **Finish** to complete the operation.

- 6 In the Create File System screen that appears next, select the **Format this volume** check box if you want to format the partition using NTFS, ReFS, FAT, or FAT32.

These options offer the following:

FAT (File Allocation Table)	A file system used by MS-DOS, Windows 3.x, and Windows 95/98. Later versions of Windows also can use the FAT file system. The operating system maintains a table to keep track of the status of various segments of disk space used for file storage.
FAT32 (File Allocation Table)	An enhanced implementation of the FAT file system. FAT32 uses smaller cluster sizes; therefore, disk space is used more efficiently, disk performance is increased, and larger drives are supported.
NTFS (NT File System)	An advanced file system designed for use specifically within Microsoft Windows operating systems. Use this format if you want to use file and folder compression. Note that NTFS cannot be used by MS-DOS, Windows 3.x, or Windows 95/98.
ReFS (Resilient File System)	An advanced file system that is built on the foundation of NTFS. ReFS is supported on Windows Server 2012 and 2012 R2 operating systems only.

- 7 Select an allocation unit size in bytes, if you want to use a size other than the default. Default settings are recommended for general use.

- 8 To name the partition, type a label in the "File system label" entry box.

Give the partition a label that is helpful to you in identifying it.

(Optional) Sselect the following formatting options:

Quick format	Use this method to skip scanning for bad sectors in the partition during formatting.
Enable file and folder compression	Use this method to compress the folders and files that are added to the volume. This option is available only if the partition is formatted with NTFS.

Click **Next** to continue.

The final screen displays the properties of the partition.

- 9 Click **Finish** in the final screen.

You can view the percentage of completion for the formatting process in the bottom pane of the console by selecting the **Tasks** tab.

New logical drives

An extended partition can be subdivided into as many as 32 logical drives. You can use all or part of the free space in an extended partition when creating logical drives. An extended partition or a logical drive cannot span multiple disks.

Note: You cannot mark a logical drive as active. An active volume or partition is the volume or partition from which the computer starts up.

To create a new logical drive

- 1 Right-click an extended partition from the left panel tree view, or click the **Disk View** tab in the right pane and right-click the extended partition in that view.
- 2 In the context menu that appears, click **New Logical Drive**.
The initial screen of the Create Partition wizard appears. Click **Next**.
- 3 Click the **Free Space** check box to select the free space to use for that logical drive. Click **Next**.
- 4 Click **Next** to move through the wizard's Select disk and region screen.
The Select Partition Type screen appears. The **Logical drive** radio button is automatically selected, and the screen displays the largest logical drive size that can be created in the extended partition.
- 5 If desired, change the size of the disk space to use, or accept it as shown.

- 6 Type the size of the logical drive that you want to create, and then click **Next**.
The Assign Drive Letter screen appears.
- 7 Choose a drive letter, mount point, or neither (a drive letter can be assigned later). Click **Next** to continue.
The Create File System screen appears.
- 8 Select the file system type and options, and then click **Next** to continue.
- 9 Review your selections and click **Finish** to complete the process.

Mark a partition as active

The Mark Partition Active command lets you designate a basic primary partition as active. An active partition is the partition from which the computer starts up. The system partition or volume must be a primary partition that has been marked as active for startup purposes and must be located on a disk that the computer accesses when starting up the system. There can be only one active system partition at a time, which displays as Active in the status field.

The Mark Partition Active command can be used only on a basic primary partition, not on a dynamic volume.

The names commonly used for partitions or volumes containing the startup and operating system files are the system volume and boot volume respectively. The system volume refers to the partition or volume containing the startup files. The boot volume is a partition or volume that contains the operating system files. The boot volume can be (but does not have to be) the same as the system volume. The boot volume also cannot be part of a spanned, striped, or RAID-5 volume.

To mark a basic volume as the active volume

- 1 Select the primary partition that contains the startup files for the operating system that you want to activate.
- 2 Right-click to display the partition's context menu, and click **Mark Partition Active** on that menu.
A message appears, advising you that the partition is started when you restart your computer.
- 3 Click **Yes** to proceed.

Change a basic disk to dynamic

To change a basic disk to dynamic, you need to create a dynamic disk group first with the New Dynamic disk group or Add Disk to Dynamic disk group command. When you do create a dynamic disk group, you can specify basic disks for the group

that become dynamic through the command. If you want to add more disks later, you use the Add Disk to Dynamic disk group command.

See [“About creating dynamic disk groups”](#) on page 43.

See [“Add a disk to a dynamic disk group”](#) on page 99.

Troubleshoot basic disks and volumes

Topics in this section are helpful in troubleshooting basic disks and volumes

See [“Bring a basic disk back to an online state”](#) on page 530.

See [“Bring a basic volume back to a healthy state”](#) on page 532.

See [“Starting and stopping the Storage Foundation Service”](#) on page 543.

Automatic discovery of SSD devices and manual classification as SSD

SFW has the ability to recognize solid-state devices (SSD) used in storage. This would help in identifying them in optimizing IO performance.

Disk media types

The following are the disk media types recognized by SFW:

- **HDD**: Stands for a hard disk drive which is used for storing and retrieving primary computer data. By default, all disks, unless auto-discovered as SSD, are recognized as HDD media type.
- **SSD**: Stands for a solid-state drive, also known as solid-state disk or a flash drive which is a data storage device used either for storage administration or as a caching device.

To manage environments containing both SSD and HDD storage devices, identify and classify them using the Volume Manager.

With this feature, you can set, change and display the media type information of the SSD disks. Once the disk is classified as SSD it is used in all space allocation operations, across nodes, restarts and failover, provided it is part of an SFW dynamic disk group.

See [“Volume Manager space allocation is SSD aware”](#) on page 144.

Alternatively, you can use SSDs in a storage environment as caching devices using the SmartIO feature.

See [“About SmartIO”](#) on page 382.

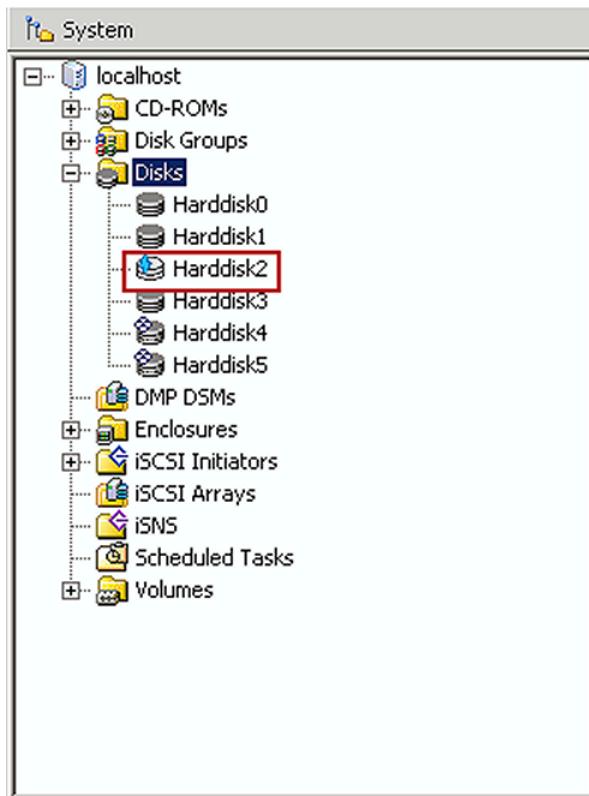
Supported Solid-State Devices

The following Solid-State Devices (SSDs) are currently supported:

- Fusion-IO
- Intel NVMe
- Violin Memory

Icon for SSD

The VEA GUI shows the following icon when it recognizes an SSD device:



Enclosure and VDID for automatically discovered On-Host Fusion-IO disks

- For auto discovered On-host Fusion-IO disks:
ARRAY_MEDIA_TYPE attribute of disk is set to 'SSD'

ARRAY_TRANSPORT is set to 'PCI'

Disk Connect Type is set as 'Internal'.

- For auto discovered On-host Fusion-IO disks, the disk VDID is as follows:
Syntax: VendorId_ProductId_LunSerNo
VendorId: "Fusion"
ProductId: "IODrive 160GB"
LunSerNo: "40279"
Disk VDID: "Fusion_IODrive 160GB_40279"
- Auto discovered On-host Fusion-IO disks, the enclosure name is as follows:
Syntax: VendorId_ProductId
Example for Fusion-IO IODrive Duo SSD storage PCI card
VendorId: "Fusion"
ProductId: "IODrive 160GB"
Enclosure VDID: "Fusion_IODrive 160GB"

Enclosure and VDID for automatically discovered On-Host Intel disks

- For auto discovered On-host Intel disks:
ARRAY_MEDIA_TYPE attribute of disk is set to 'SSD'
ARRAY_TRANSPORT is set to 'PCIe'
Disk Connect Type is set as 'Internal'.
- For auto discovered On-host Intel disks, the disk VDID is as follows:
Syntax: VendorId_ProductId_CabSerNo_LunSerNo
VendorId: "INTEL(R)"
ProductId: "SSD 910 200GB"
CabSerNo: "INTEL"
LunSerNo: "5000CCA013008CE7"
Disk VDID: "INTEL(R)_SSD 910 200GB_INTEL_5000CCA013008CE7"
- Auto discovered On-host Intel disks, the enclosure name is as follows:
Syntax: VendorId_ProductId_CabSerNo
VendorId: "INTEL(R)"
ProductId: "SSD 910 200GB"
CabSerNo: "INTEL"
Enclosure: "INTEL(R)_SSD 910 200GB_INTEL"

Enclosure and VDID for automatically discovered Violin disks

- For auto discovered Violin disks:
ARRAY_MEDIA_TYPE attribute of disk is set to 'SSD'

Disk Connect Type is set as 'External'

- For auto discovered Violin disks, the disk VDID is as follows:
Syntax: VendorId_ProductId_CabSerNo_LunSerNo
VendorId: "VIOLIN"
ProductId: "SAN ARRAY ALUA"
CabSerNo: "veritas-cert"
LunSerNo: "34DC241D7C1EDCB6"
Disk VDID: "VIOLIN_SAN ARRAY ALUA_veritas-cert_34DC241D7C1EDCB6"
- Auto discovered Violin disks, the enclosure name is as follows:
Syntax: VendorId_ProductId_CabSerNo
VendorId: "VIOLIN" ProductId: "SAN ARRAY ALUA"
CabSerNo: "veritas-cert"
Enclosure: "VIOLIN_SAN ARRAY ALUA_veritas-cert"

Classifying disks as SSD

An SSD can be an internal disk or external disk. You can manually classify any disk as SSD or HDD, change, and display the media type information.

You can set the media type and connection type for a disk using the following `vxdisk setmedia` command:

```
Command: vxdisk setmedia type=<HDD|SSD> [-x] <diskname | p#c#t#l#>
```

Here is an example of the `setmedia type` command showing how to set the media type of a disk:

```
C:\> vxdisk setmedia type=ssd -x harddisk4
```

```
Successfully set media type of specified disks.
```

Note: If you do not specify the disk connection type as external, then by default, that particular disk is marked as internal.

You can get the media type and connection type information of the disk using the following `vxdisk` command:

```
Command: vxdisk diskinfo -e <diskname | p#c#t#l#>
```

Here is an example of the `vxdisk diskinfo` command displaying the new state:

```
C:\>vxdisk diskinfo -e Harddisk6  
Disk information  
Device Name           : Harddisk6
```

```
...  
Array Media Type : SSD  
...  
Disk Connection Type: External  
...
```

You can get the plex media type information for the given volume using the following `vxvol volinfo` command:

```
Command: vxvol -v volinfo <driveletter>
```

Here is an example of the `vxvol` command displaying the media type of plexes:

```
C:\> vxvol -v volinfo h:  
Volume information  
Name           : v3 (H:)  
DeviceName     : \Device\HarddiskDmVolumes\SSDg\v3  
DriveLetter    : H:  
Disk Group     : SSDg  
...  
Names of the plexes are...  
Plex 0         : v3-01  
...  
Media Type : SSD
```

You can get the plex media type information for all the existing plexes using the following `vxprint` command:

```
Command: vxprint -l -p
```

Here is an example of the `vxprint` command displaying the media type of plexes:

```
C:\> vxprint -l -p  
Diskgroup = BasicGroup  
Diskgroup = SSDg  
Plex       : v3-01  
info      : len=102774784 KB  
assoc     : vol=v3 sd=Disk1-02  
flags     : ncolumn=1 StripeWidth=0  
GUID      : 72bdc94b-ded8-4f2f-9f30-2cdd8cabc764  
type      : layout=Simple  
...  
Media Type : SSD  
...
```

A plex can have the following four media types:

- **SSD:** All subdisks of the plex reside on SSD disks.
- **HDD:** All subdisks of the plex reside on HDD disks.
- **Mixed:** The subdisks of the plex reside on SSD and HDD disks.
- **Unknown:** If some disks on the plex are not available.

Limitations for classifying SSD devices

The limitations for classifying SSD devices are as follows:

- No GUI support is available for manually classifying a disk as SSD.
- Classification is allowed only if the disk is part of an SFW dynamic disk group.
- Classification of SFW dynamic disks is allowed only if it is in read-write imported state.

See [“Volume Manager space allocation is SSD aware”](#) on page 144.

See [“Setting an SSD plex as a preferred plex to improve I/O performance”](#) on page 145.

Volume Manager space allocation is SSD aware

Volume Manager space allocation operations are now media type aware and currently supports the media types SSD and HDD.

The following space allocation operations are SSD aware:

- **For volumes:** Volume Create, Add Mirror, Volume Grow, Add Logs, Snap Prepare, and storage migration.
- **For disks:** Disk Replace, Disk Evacuate, and Subdisk Move.

Any new plex cannot span across media types. The complete plex would either be on an SSD or HDD.

Media type restrictions are not considered for site separated volumes.

If the existing plex is media confined, the Volume Grow operation respects the media type boundary. Otherwise, the plex may be extended to either HDDs or SSDs.

For Subdisk Move and related operations, such as Disk Evacuate, and volume migration without changing the layout, a subdisk can be moved only to a target disk of the same media type as the source. However, if you opt to change the layout of the volume during migration, then the media type restriction is not considered while choosing the target.

The Array migration operation fails if HDD target disks are assigned to the SSD source disks. In such scenarios, you may change the media type of the source disks using `vxdisk setmedia` and perform Array migration.

Further, during hot relocation, the media type restriction is not considered.

See [“Disk media types”](#) on page 139.

See [“Setting an SSD plex as a preferred plex to improve I/O performance”](#) on page 145.

See [“Automatic discovery of SSD devices and manual classification as SSD”](#) on page 139.

Setting an SSD plex as a preferred plex to improve I/O performance

A volume may contain plexes from media types HDD and SSD.

The Read performance of an SSD plex is better than an HDD plex. Therefore, the SSD plex is set as the preferred plex with default Read policy based on the following conditions:

- If you have already set a preferred plex then the above mentioned changes do not take place.
- No SSD plex is set as a preferred plex if all plexes of a volume lie on SSD disks.
- The preferred plex setting on an SSD plex can be controlled using the following registry value: `SetPreferredPlexonSSD`

Run the command `Regedit` to open the registry editor.

Locate the `SetPreferredPlexonSSD` in the registry under the following key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager.
```

Automatic selection of an SSD plex as a preferred plex can be turned off by setting the above DWORD registry value to '0'.

- If a non-empty disk is manually tagged, then the preferred plex settings of the existing volumes on that disk is not updated.

See [“Volume Manager space allocation is SSD aware”](#) on page 144.

See [“Disk media types”](#) on page 139.

See [“Automatic discovery of SSD devices and manual classification as SSD”](#) on page 139.

Dealing with disk groups

This chapter includes the following topics:

- [Disk groups overview](#)
- [Delete a dynamic disk group](#)
- [Upgrading the dynamic disk group version](#)
- [Converting a Microsoft Disk Management Disk Group](#)
- [Importing a dynamic disk group to a cluster disk group](#)
- [Rename a dynamic disk group](#)
- [Detaching and attaching dynamic disks](#)
- [Importing and deporting dynamic disk groups](#)
- [Partitioned shared storage with private dynamic disk group protection](#)
- [Dynamic disk group properties](#)
- [Troubleshooting problems with dynamic disk groups](#)

Disk groups overview

Storage Foundation for Windows organizes disks into disk groups. Disk groups provide a way of organizing disks and simplifying storage management for systems with large numbers of disks. They also let you move disks between computers so that you can easily transfer the storage between computers.

See [“Basic and dynamic disk groups”](#) on page 147.

See [“Moving dynamic disk groups between computers”](#) on page 147.

See [“Cluster disks and cluster dynamic disk groups”](#) on page 148.

See [“Private dynamic disk group protection”](#) on page 148.

See [“About creating dynamic disk groups”](#) on page 43.

See [“About Dynamic Disk Group Split and Join”](#) on page 273.

Basic and dynamic disk groups

Storage Foundation for Windows assigns one group for all basic disks, called the Basic disk group. Dynamic disks, however, can have multiple disk groups, which are called dynamic disk groups. Disks within a dynamic disk group share a common configuration. Dynamic volumes are created within a dynamic disk group and are restricted to using disks within that group.

When SFW is run for the first time on a server, all disks are in the Basic disk group. To upgrade basic disks to dynamic, one or more dynamic disk groups have to be created with the New Dynamic Disk Group command.

Moving dynamic disk groups between computers

You can move a dynamic disk group between computers by deporting it on one computer, moving the disk or disks that contain the dynamic disk group to another computer, rescanning on the new computer, and then importing the dynamic disk group. You can use the volumes from the imported dynamic disk group without having to restart if you have hot-swappable disks.

Before moving the disks in a dynamic disk group, make sure that the disks are online and the volumes are healthy. You should move all disks that are part of the dynamic disk group; otherwise, the volumes have Degraded status when the disks are moved to the other computer.

Primary and secondary dynamic disk groups

Storage Foundation for Windows makes a distinction between primary and secondary dynamic disk groups. In the user interface, the primary or secondary designation is shown in parentheses after each dynamic disk group name.

In Storage Foundation for Windows and all versions of Volume Manager for Windows, only one primary dynamic disk group can exist on a single host computer. Additional groups that are created or imported on that computer are secondary dynamic disk groups.

Note: A primary disk group upgraded from Microsoft Disk Management with SFW running on Windows Server operating systems always becomes a secondary dynamic disk group.

You might wonder what happens in later releases of Volume Manager or SFW if you have a primary dynamic disk group that was upgraded from Microsoft Disk Management or earlier versions of Volume Manager and that has no boot or system disk and then you create another dynamic disk group that has the computer's boot or system disk. Once this second disk group is created, it becomes the primary dynamic disk group, and the other disk group that was upgraded from the earlier software becomes a secondary dynamic disk group.

The primary dynamic disk group that contains the computer's system or boot disk cannot be deported because doing so would make the computer unbootable.

Note: A primary dynamic disk group should not be renamed.

The concept of importing dynamic disk groups is also used in Storage Foundation for Windows in connection with a single computer. If a computer is restarted, SFW makes available or "auto-imports" the primary dynamic disk group and any secondary dynamic disk groups that previously existed on the computer before the restart, providing these groups have not been deported to another computer.

Cluster disks and cluster dynamic disk groups

Storage Foundation (SFW) has a special category of disk groups for disks involved in supporting Cluster Server (VCS) and Microsoft Failover Clustering.

In order for SFW to work with VCS or Microsoft Failover Clustering, you create a cluster dynamic disk group for the disks that are part of the cluster. To create a cluster dynamic disk group, you use the New Dynamic Disk Group command and click the **Create Cluster Dynamic Disk Group** check box.

More information about Microsoft Failover Clustering is available here at:

See "[Create a cluster dynamic disk group](#)" on page 500.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Private dynamic disk group protection

This feature lets you partition shared storage on a SAN or shared array for exclusive ownership by a single machine. The storage is partitioned by using a secondary dynamic disk group. Private dynamic disk group protection provides hardware locking to the disks in the secondary dynamic disk group through a SCSI reservation thread.

See [“Partitioned shared storage with private dynamic disk group protection”](#) on page 162.

Note: Microsoft Disk Management Disk Groups do not support private dynamic disk group protection.

Delete a dynamic disk group

Deleting a dynamic disk group can be done two ways.

The recommended procedure is a two-part process and is described as follows:

- First, all the volumes in the dynamic disk group must be moved or deleted by using the Delete Volume command, then all the disks in the dynamic disk group must be removed from the disk group by using the Remove Disk from Dynamic Disk Group command. Once you remove all the disks from a dynamic disk group, the dynamic disk group is automatically deleted.
- The alternative procedure is to use the Destroy Dynamic Disk Group command, which permanently removes the disk group and all its volumes in a one-step process.

Note: All data in a volume is lost when you delete the volume. Make sure that the information has been transferred elsewhere or that it is no longer needed.

Recommended method for deleting a disk group

Perform the following steps to delete a disk group.

To delete a dynamic disk group by using the recommended method

- 1 Delete each volume by right-clicking on the volume in the tree view or in the right-pane Disk View. Select **Delete Volume** from the context menu.
- 2 You are prompted for verification. Click **Yes** to delete the volume.
- 3 Repeat steps 1 and 2 until all volumes in the dynamic disk group have been deleted.
- 4 Right-click on the dynamic disk group to be deleted. Select **Remove Disk from Dynamic Disk Group** from the context menu.

- 5 In the **Remove Disk** window that appears, use the **Add** or **Add All** option to move all the disks in the dynamic disk group to the **Selected Disks** column in the right pane of the window.
- 6 Click **OK**. The dynamic disk group is removed from the tree view automatically.

Alternative method for deleting a disk group

The alternative method for deleting a disk group is to use the Destroy Dynamic Disk Group command. Because this command permanently removes the disk group and all its volumes, please use caution in implementing this command.

Note: Microsoft Disk Management Disk Groups do not support the Destroy Dynamic Disk Group command.

Warning: This command permanently removes all volumes and disks within the disk group.

To delete a dynamic disk group with the alternative method

- 1 Right-click the storage object representing the server being managed (the computer icon at the top of the tree for the server being managed).
- 2 Select **Destroy Dynamic Disk Group** from the context menu.
- 3 Enter the name of the dynamic disk group you want to delete in the dialog box that appears.
- 4 Click **OK**.
A confirmation box appears asking if you want to destroy the disk group.
- 5 Click **Yes** to delete the dynamic disk group or **No** to cancel the operation.
After the dynamic disk group is destroyed, the disks revert to basic disks and become part of the Basic disk group.

Upgrading the dynamic disk group version

This task is applicable only if you have upgraded the product from a version prior to 6.1.

The dynamic disk group version is not upgraded automatically after you upgrade SFW or SFW HA. If a service group in the previous configuration contains a dynamic disk group, you must upgrade its version after you complete the SFW or SFW HA upgrade.

Please note that the following features and operations are introduced with the new disk group version and, therefore, will not be supported if the disk group version is not upgraded:

- The snapshot feature (prepare, snap shot, snap back, snap abort, and snap clear operations)
- The dirty region logging (DRL) feature (add a log and remove a log operations)
- Fast Resync and Fast File Resync features for volumes (add a log and remove a log operations)
- Import disk group (import a legacy disk group as a shared disk group operation)
- SSD caching

Note: When a legacy disk group with dirty region logging (DRL) logs in its volumes is upgraded, the DRL logs are deleted on upgrade. You need to add the DRL logs manually after the upgrade.

Note: Once a disk group version is upgraded, it cannot be changed back to an earlier disk group version and you can not import it on another server that is running any earlier versions of SFW.

Note: Microsoft Disk Management Disk Groups do not support upgrading the disk group version. However, On Windows Server 2008 R2 and 2010, you can first convert the disk group to an SFW dynamic disk group, and then upgrade the SFW dynamic disk group.

Note: After upgrading the disk group to the latest version and importing a disk group as a cluster-shared disk group (CSDG), you need to manually create the Volume Manager Shared Volume resource for each volume in a disk group.

See [“Manually creating a Volume Manager Shared Volume resource”](#) on page 431.

You can upgrade the dynamic disk group using the Veritas Enterprise Administrator (VEA) console or the CLI.

To upgrade the dynamic disk group version from the command line:

- 1 To upgrade the disk group using CLI, run the following command. You must run this command for each dynamic disk group separately.

```
vxdg -gDynamicDiskGroupName [-T version] upgrade
```

Where,

DynamicDiskGroupName = Name of the dynamic disk group

version = Target version of the dynamic disk group

- 2 After the disk group version is upgraded successfully, the disk group version should be 171. Run the following command to verify the upgraded version:

```
vxdg -gDynamicDiskGroupName dginfo
```

To upgrade the dynamic disk group version using VEA

- 1 From the VEA console, right-click the disk group. You must perform these steps for each dynamic disk group separately.
- 2 Select **Upgrade Dynamic Disk Group Version**.

A notification indicating that the disk group version has been upgraded is seen in the alert logs. After the disk group version is upgraded successfully, the disk group version should be 171. You can verify the disk group properties to confirm the upgraded version. After the disk group version is upgraded successfully, the disk group version should be 171.

Converting a Microsoft Disk Management Disk Group

On Windows Server operating systems, converting a Microsoft Disk Management disk group to a SFW dynamic disk group can be done by using the Convert to SFW Dynamic Disk Group command.

Note: Converting a Microsoft Disk Management Disk Group that was created using iSCSI disks to an SFW dynamic disk group is currently not supported.

To convert a Microsoft Disk Management disk group to an SFW dynamic disk group

- 1 In the tree view, right-click the Microsoft Disk Management disk group and select the command **Convert to SFW Dynamic Disk Group** from the context menu.
- 2 The Convert to SFW Dynamic Disk Group confirmation box appears.
Review the steps of the procedure and click **Yes** to begin the conversion.
- 3 Restart the server.
- 4 When the server completes the restart, launch SFW and import the converted disk group.

Note: You must specify to clear the host id for the import operation.

Note: The default name of the disk group may not be user-friendly. You can change the name of the disk group in the dialog box of the import command.

- 5 Upgrade the imported disk group, if appropriate.
The imported disk group is a Microsoft Disk Management compatible disk group managed by SFW. If advanced disk group features are needed, upgrade the disk group version with the **Upgrade Dynamic Disk Group Version** command.

See [“Upgrading the dynamic disk group version”](#) on page 150.

Importing a dynamic disk group to a cluster disk group

Importing a dynamic disk group as a cluster disk group can be done by using the Import as cluster disk group option of the Import Dynamic Disk Group command. This conversion to a cluster disk group requires that the dynamic disk group be connected to a shared bus.

See [“Import a dynamic disk group”](#) on page 159.

Note: Microsoft Disk Management Disk Groups do not support the Import Dynamic Disk Group command.

To import a dynamic disk group to a cluster disk group

- 1 In the tree view, right-click the disk name in the dynamic disk group or the dynamic disk group name in the tree view and select Import Dynamic Disk Group.
- 2 Select the Import as cluster disk group option to import the dynamic disk group as a cluster disk group.

Rename a dynamic disk group

Perform the following steps to rename a disk group.

Note: You should not rename a primary dynamic disk group if it contains the boot or system volume.

Note: Microsoft Disk Management Disk Groups do not support the Rename Dynamic Disk Group command.

To rename a dynamic disk group

- 1 In the tree view, right-click the disk group and select the command **Rename Dynamic Disk Group** from the context menu.
- 2 The Rename Dynamic Disk Group dialog box comes up. Type the new name in the **New Name** text entry box, and click **OK**.

Note: The dynamic disk group name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

Detaching and attaching dynamic disks

Storage Foundation for Windows supports detaching and attaching disks of dynamic disk groups. This feature lets you detach dynamic disks for maintenance or other reasons and then attach the dynamic disks when they become available for use. Detached disks are no longer accessible and cannot be operated on until they are reattached.

Detaching and attaching disks using the VEA GUI is described in this section.

You can also detach and attach a disk using the `vxdisk` CLI command.

See “`vxdisk`” on page 624.

Detaching dynamic disks

The Detach Disk command stops disk access and disconnects the disks from SFW control.

Note the following:

- The detach disk operation can only be performed on the disks that are in the attached state.
- You may detach the disks that may or may not contain any data.
- This operation can be used on all the disks in a dynamic disk group, however at least one disk in the dynamic disk group must remain in attach state to maintain the last active configuration or log copy.
- Detaching disks of dynamic disk groups is not recommended for the disks that contain system or boot volumes, or the disks that contain a page file or quorum resource.

Note: This operation fails the volume if all the plexes of a redundant volume are contained only on the disks specified for the detach operation.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Detach Disk command.

To detach disks

- 1 Stop all processes using the volumes in the dynamic disk group, including all I/O to files.

If you do not stop the processes using the volumes, there is the possibility that data written to internal buffers may not get written to disk, resulting in data loss. It is also recommended that you back up all data on your volumes before you detach disks.

- 2 Make sure the status of the volumes on the disks is Healthy and the disks are online.

If the status is not Healthy, you should repair the volumes before you detach the disks.

- 3 Select the disks in a dynamic disk group or the individual disks, and select **Detach Disk** from the menu.

A warning message asks if you want to detach the disks.

- 4 Click **Yes** to detach the disk.

When disks are detached, their redundant volumes are shown as degraded. For non-redundant volumes, the disks are shown as failed.

If a disk is enabled for Hot Relocation and is detached, its subdisks are not hot relocated. Also subdisks cannot be hot relocated from other disks to a disk that has been detached.

- 5 Remove the disks from the system.

You may need to prepare disks for removal in some cases. See your hardware manual for information.

Attaching dynamic disks

The Attach Disk command connects the disks and enables access to the disks.

Note the following:

- The attach disk operation can only be performed on the disks that are in the detach state.
- You may attach disks that may or may not contain any data.
- When attaching a disk and bringing it online, the disk's volumes should be reactivated and any redundant volumes should be synchronized automatically.
- When an Attach Disk command is issued for the disks, the disks come back online and volumes return to the state they were in at the time of the detach (assuming that all the disks in the disk group are present).

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Attach Disk command.

To attach disks

- 1 Select the disks in a dynamic disk group or the individual disks, and select **Attach Disk** from the menu.

- 2 Click **Yes** to attach the disks.

The **Attach Disk** command automatically resynchronizes the plexes of redundant volumes that are contained on the specified disks. However, non-redundant volumes must be reactivated manually.

Importing and deporting dynamic disk groups

Storage Foundation for Windows supports deporting and importing of dynamic disk groups. This feature lets you move disks belonging to a dynamic disk group between computers and make their storage available for use. If you have hot-swappable disks, it is possible to move disks between systems without a system shutdown and restart.

Note: You must have administrator privileges on the source and destination computers to be able to move Storage Foundation for Windows storage between computers.

Deport a dynamic disk group

The `Deport Dynamic Disk Group` command stops access to disks. If you move hot-swappable disks between computers, it is important to use the `Deport Dynamic Disk Group` command. Using the command ensures that the data has been preserved in a clean state before you move the disks to the other computer. The `Deport Dynamic Disk Group` command also clears the host ID of the computer on which the disk group is located, so that it can be imported on another computer.

Warning: Though it is possible to deport dynamic disk groups containing the volumes that are in a Degraded state, this is not recommended. If a disk failed during transport between systems or during the import, it is possible that data would be lost for volumes running in the Degraded state.

Note: Microsoft Disk Management Disk Groups do not support the `Deport Dynamic Disk Group` command.

Note: Storage Foundation for Windows does not let you deport a cluster disk group that is a Microsoft Failover Clustering disk group resource.

To deport a dynamic disk group

- 1 Stop all processes using the volumes in the dynamic disk group, including all I/O to files.

If you do not stop volume processes, there is the possibility that data written to internal buffers may not get written to disk, resulting in data loss. It is also recommended that you back up all data on your volumes before you move disks to another computer.

- 2 Make sure the status of the volumes on the disks is Healthy and the disks are online.

If the status is not Healthy, you should repair the volumes before you move the disks.

- 3 Select a disk in the dynamic disk group or the dynamic disk group name, and select **Deport Dynamic Disk Group** from the menu that appears.

A warning message asks if you are sure that you want to deport the disk group.

- 4 Click **Yes** to deport the dynamic disk group or **No** if you decide not to deport the dynamic disk group.

The process of deporting a dynamic disk group puts the contained disks in the Offline state and all volumes in the Stopped state. This applies only while the dynamic disk group is deported. Once an **Import Dynamic Disk Group** command is issued, disks come back online and volumes return to the state they were in at the time of the deport (assuming that all the disks in the disk group are present).

- 5 Remove the disks from the current system.

If the disks are hot swappable, you can move disks between systems without performing a system shutdown. If the drives are NOT hot swappable, it is necessary to turn off the computer to remove or install the drives.

You may need to prepare disks for removal in some cases. See your hardware manual for information.

Be sure to move all the disks that are members of the dynamic disk group at the same time.

Normally, you need to move all the disks in the dynamic disk group to the other computer. The only exception would be if you move the disks temporarily to another computer to copy one or more of the volumes that resided on some of the disks. In that case, when you import the disks to the other computer, the volumes that are not complete because of missing disks show a Missing status. If you reactivate the disks and the volumes, the true status of the volumes is displayed, such as Failed or Degraded. Once you move the disks back to the original computer, rescan, and reimport the dynamic disk group, all the volumes in the dynamic disk group should come back in a Healthy state. You may need to use the **Merge Foreign Disk** command.

See [“Merge foreign disk command”](#) on page 538.

Import a dynamic disk group

Perform the following steps to import a disk group.

Note: Normally, you cannot import a dynamic disk group on a computer without having deported it on the computer where the disk group was located previously. However, the Import Dynamic Disk Group command has an option to clear the host ID of the previous computer, so that you can do an import operation. This option should be used with caution.

Note: While importing a cluster disk group as a cluster-shared disk group (CSDG), you must select the “Clear Host ID” check box.

Note: After upgrading the disk group to the latest version and importing a disk group as a cluster-shared disk group (CSDG), you need to manually create the Volume Manager Shared Volume resource for each volume in a disk group. To do this, See [“Manually creating a Volume Manager Shared Volume resource”](#) on page 431.

Note: Microsoft Disk Management Disk Groups do not support the Import Dynamic Disk Group command.

Once you have deported a dynamic disk group and disconnected the disks from the original machine, do the following to import the dynamic disk group

To import a dynamic disk group:

- 1 Reinstall the hardware in the destination computer.
- 2 If the disks are hot swappable, open the VEA console on the destination computer and issue a **Rescan** command. If the system had to be shut down to install the disks, it will have done a rescan automatically during the restart.

Once the rescan completes or the system restarts, the dynamic disk group and disks should be listed; but the dynamic disk group is in a Deported state, and the disks are in the Offline or Foreign state.

Occasionally, you may have to rescan a second time to get the disks showing as Offline or Foreign, particularly if you do the first rescan very quickly after installing the disks on the second computer system.

- 3 Right-click on a disk name in the dynamic disk group or the dynamic disk group name in the tree view, and select **Import Dynamic Disk Group** from the menu that appears.

- 4 If you want to change the dynamic disk group's name, type in a new name for the group in the screen that comes up.

Note: The dynamic disk group name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

If the disk group you import has a name longer than 18 characters (as was possible in earlier versions of Volume Manager for Windows), you may be asked to shorten the disk group name. It is possible that you will be asked to shorten the name to fewer than 18 characters if the disk group's volumes also have very long names.

Do not rename a dynamic disk group that contains a system or boot volume, otherwise an error can occur.

See [“Error that the boot device is inaccessible, bugcheck 7B”](#) on page 554.

On this screen, you may also check **Clear host ID** to clear the host id of the previous computer that had the dynamic disk group.

This screen also lets you import the dynamic disk group as a cluster disk group. Select **Import as cluster disk group** to import the dynamic disk group as a cluster disk group. Importing the dynamic disk group to a cluster disk group requires that the dynamic disk group be connected to a shared bus.

The **Import as cluster disk group** option is already selected if you import a cluster disk group. Select **Import as dynamic disk group** to import the cluster disk group as a dynamic disk group.

In addition, you can specify that the dynamic disk group be imported as a SAN disk group. This option is already selected if you import a SAN disk group. Select **Import as SAN disk group** to import the dynamic disk group as a SAN disk group.

- 5 If you have not done a deport command on the computer where the disk group was located previously, click the check box to clear the host ID of the other computer.
- 6 Click **OK** to import the dynamic disk group.

If the volumes in the dynamic disk group come up as Failed or Degraded, you need to repair the volumes.

See [“Repair a volume with degraded data after moving disks between computers”](#) on page 534.

Partitioned shared storage with private dynamic disk group protection

Partitioned shared storage with private dynamic disk group protection allows administrators to partition storage and assign exclusive ownership in a SAN or shared disk array. In early releases of Volume Manager for Windows, partitioned shared storage was implemented through cluster disk groups without the use of a cluster application. In later releases of Volume Manager and Storage Foundation for Windows, the feature of private dynamic disk group protection makes it possible to partition shared storage with a regular dynamic disk group.

Private dynamic disk group protection uses hardware locking techniques to protect secondary dynamic disk groups located on shared storage from access by other hosts connected to the shared storage pool. The hardware locking is implemented by using a SCSI reservation thread to maintain a current reservation for each disk in a protected group.

In later releases of Volume Manager and SFW, partitioned shared storage is available for secondary dynamic disk groups on shared storage. Private dynamic disk group protection can be assigned when a secondary dynamic disk group is created, or it can be added to an existing secondary dynamic disk group. The protection is in place as long as the disk group is imported; and since secondary dynamic disk groups are auto-imported, protected disk groups come online at system startup with a SCSI reservation thread started automatically.

If disks are being added to a dynamic disk group that has a SCSI reservation, they get a reservation automatically. However, if a new disk group is created, you must choose to add private group protection to the disk group.

Private dynamic disk group protection can be removed if the disk group no longer needs protection or is reassigned to another Storage Foundation for Windows host.

The group protection is automatically removed and needs to be added again in the following cases:

- Dynamic disk group is deported and moved to another host.
- The product is uninstalled and then reinstalled.

Note: Microsoft Disk Management Disk Groups do not support private dynamic disk group protection.

See [“Create a new dynamic disk group with private dynamic disk group protection”](#) on page 163.

See [“Add private dynamic disk group protection to an existing dynamic disk group”](#) on page 164.

See [“Remove private dynamic disk group protection from a dynamic disk group”](#) on page 165.

Create a new dynamic disk group with private dynamic disk group protection

The following steps describe how to use the GUI to partition shared storage using private dynamic disk group protection.

Note: Microsoft Disk Management Disk Groups do not support private dynamic disk group protection.

To use the GUI to partition shared storage using private dynamic disk group protection

- 1 Identify the set of shared basic disks you want to assign to a particular Storage Foundation for Windows server.
- 2 Use the VEA console to connect to Storage Foundation for Windows on that computer.
- 3 Select one of the basic disks you want to upgrade, and select **New Dynamic Disk Group** from the context menu.
- 4 The initial screen of the New Dynamic Disk Group wizard appears. Click **Next** to continue.
- 5 The screen for entering the disk group attributes comes up. Do the following:
 - Type a name for the new group in the **Group Name** entry box.
 - Click the **Add private group protection** check box.
 - Move each of the basic disks that you want to include in the new disk group from the left pane of the wizard to the right pane by using the **Add** button.
 - Click **Next** to continue.
- 6 Click **Next** to confirm the group of disks to be included in the new dynamic disk group.
- 7 Select **Finish** on the next screen to create the new group.

The icon for the new disk group includes a lock icon superimposed on its lower right quadrant. The disk group Properties display shows that the group is of type "Dynamic(Secondary, Private)."

Using the vxdbg init command

It is possible to perform the operation described above by using the following command-line command:

```
vxdbg init -g<NewDynamicGroupName> -R <DiskName> ...
```

where *<DiskName>* is of the form `harddisk0`, `harddisk1`, etc.

Disks added to the group, either through the Add Disk to Dynamic Disk Group wizard or by a Join Dynamic Disk Group operation, are automatically protected as soon as the add operation is complete. Disks that are removed from the group, using either the Remove Disk from Dynamic Disk Group or Split Dynamic Disk Group operations, are no longer protected.

Add private dynamic disk group protection to an existing dynamic disk group

Perform the following steps to add private dynamic disk group protection to an existing disk group.

Note: Microsoft Disk Management Disk Groups do not support private dynamic disk group protection.

To add private dynamic disk group protection to an existing secondary disk group

- 1 Right-click dynamic disk group in the tree view of the VEA console.
- 2 Select **Add Private Dynamic Disk Group Protection**.
- 3 A pop-up confirmation window appears. Click **Yes**.

The disk group's icon now includes a lock icon superimposed on its lower right quadrant. The disk group properties display shows that the group is of the type Dynamic(Secondary, Private).

Using the vxdbg protect command

It is possible to perform the operation described above by using the following CLI command:

```
vxdbg -g<DynamicDiskGroupName> protect
```


Remove private dynamic disk group protection from a dynamic disk group

Perform the following steps to remove private dynamic disk group protection from a disk group.

Note: Microsoft Disk Management Disk Groups do not support private dynamic disk group protection.

To remove private dynamic disk group protection from a disk group

- 1 Right click a dynamic disk group from the left pane of the VEA console.
- 2 Select **Remove Private Dynamic Disk Group Protection**.
- 3 A pop-up confirmation window appears. Click **Yes**.

The lock icon previously superimposed on the disk group icon's lower right quadrant is removed. The disk group properties display shows that the group is of type Dynamic(Secondary).

Using the vxdbg release command

It is possible to perform the operation described above by using the following CLI command:

```
vxdbg -g<DynamicDiskGroupName> release
```

Dynamic disk group properties

If you select a dynamic disk group and then select Properties from the context menu, the Properties for that dynamic disk group displays.

The screen below is an example of the Dynamic Disk Group Properties screen. The status category is either Imported or Deported. "Current Version" refers to the disk group version that supports advanced disk group features.

The categories of a disk group's properties are described in the table below:

Table 5-1 Dynamic Disk Group Properties Categories

Category	Description
Name	The disk group name is shown at the top of the window and in the title bar.
Status	Fields are Imported and Deported.

Table 5-1 Dynamic Disk Group Properties Categories (*continued*)

Category	Description
Type	Fields are Dynamic (Primary) and Dynamic (Secondary).
Disks	Number of disks in the disk group.
Volumes	Number of volumes in the disk group.
Version	The disk group version. See the description of disk group versions below this table.
Recovery needed	Fields are True and False. Indicates whether the disk group has volumes that are degraded.
Shared	Fields are Yes and No. Indicates whether the disk group has private dynamic disk group protection. See "Partitioned shared storage with private dynamic disk group protection" on page 162.
Enabled	Fields are Yes and No. Indicates whether the disk group is functioning or not.
Size	Size of the disk group.
Free space	Free space remaining in the disk group.
Comment	User-defined comment.

Note: Microsoft Disk Management Disk Groups do not support disk group comments.

The following are the disk group versions:

Version number	Version Used in:
20	Microsoft Windows Disk Management, Volume Manager for Windows NT, Volume Manager 2.5 and 2.7 for Windows 2000
30	Volume Manager 3.0 and 3.1 for Windows 2000
40	Storage Foundation 4.0, 4.1, and 4.2
43	Storage Foundation 4.3
50	Storage Foundation 5.0 and 5.1
60	Storage Foundation 6.0, 6.0.1, and 6.0.2

Troubleshooting problems with dynamic disk groups

The following are troubleshooting topics in the Troubleshooting chapter.

See [“Repair a volume with degraded data after moving disks between computers”](#) on page 534.

See [“A disk is marked as foreign”](#) on page 545.

Fast failover in clustered environments

This chapter includes the following topics:

- [What is fast failover](#)
- [SFW changes for fast failover](#)
- [Enabling fast failover for disk groups in a VCS environment](#)
- [Enabling fast failover for disk groups in a Microsoft Failover Clustering environment](#)
- [Limitations](#)

What is fast failover

Fast failover is a new feature that improves the failover time for the storage stack configured in a clustered environment. Fast failover includes several design changes and enhancements to the core SFW components. These changes provide significant reduction in the failover time taken by storage resources during service group failovers.

The following factors determine how fast SFW fails over the storage stack:

- **Fault detection**
How fast the clustering software detects a fault
- **Fault notification**
How fast SFW is able to get the fault notification
- **Storage management and remedial actions**
How SFW manages configuration changes and the actions it takes on the storage

With fast failover, SFW addresses Fault notification and Storage management and remedial actions factors. SFW has optimized the algorithms and enhanced the internal logic used to manage disk groups and the storage configuration changes across the nodes in a cluster.

Cluster Server (VCS) addresses the Fault detection factor with the Intelligent Monitoring Framework (IMF) feature that provides instantaneous fault detection capability.

Fast failover integrates with the IMF feature to provide a significant performance improvement in SFW HA cluster environments.

SFW changes for fast failover

The following changes are introduced in SFW to support fast failover:

- New dynamic disk group state **Deported Read-Only**

SFW supports two disk group states; Imported and Deported.

When storage resources are online on a node (the active node), the disk group is imported in a Read/Write mode. In this state the disk group is accessible and writable on the active node. When storage resources are offline on a node (the passive node), the disk group is in a deported state. In this state the disk group is not accessible from the passive node.

With fast failover, a new state, **Deported Read-Only**, is added to the SFW dynamic disk group entity. The disk group status on passive nodes now displays as Deported Read-Only.

In this state, the disk group is imported in a “Read-Only” mode on the passive node. All the volumes are in an offline state and no transactions are allowed on the disk group.

During failover, SFW normally performs a complete disk group deport operation on the active node followed by a Read/Write import operation on a passive node. With fast failover, instead of performing deport and import operations, SFW now performs only a mode change for the disk group. The disk group state on the passive node is changed from Read-Only to Read/Write.

A mode change (Read-Only to Read/Write) is a much faster operation as compared to a full deport and import (Deport None to Import Read/Write) and thus results in faster disk group failovers.

The following table describes the different disk group states. The 4th column shows the new Deported, Read-Only state introduced as part of fast failover.

Disk group property	Imported, Read/Write	Deported, None	Deported, Read-Only
Can disk group configuration be modified?	Yes	No	No
Are volume devices created?	Yes	No	Yes
Are disks reserved?	Yes	No	No

- Changes to vxdbg command

The SFW vxdbg command is modified to display the new disk group state introduced for fast failover. The vxdbg list and vxdbg dginfo options now display the new Deported Read-Only state of fast failover enabled disk groups. Here's an example of the vxdbg list command displaying the new state:

```
C:\>vxdbg list
Disk groups are....
Name           :   dgl
Disk Group Type :   Dynamic (Cluster)
DiskGroup ID   :   689ef763-9c1f-4e96-807e-1b97543b3dba
Status         :   Deported
Access         :   Read-Only
```

Here's an example of the vxdbg dginfo command displaying the new state:

```
C:\>vxdbg -gSQL_DG dginfo
Disk group information for the disk group SQL_DG
Name           :   SQL_DG
DiskGroup ID   :   b1f9e838-0689-4b91-b737-cd33f60cdf4
Disk Group Type :   Dynamic (Cluster)
Private Group Protection :   No
Status         :   Deported
Access         :   Read-Only
Current Version :   Yes
Version Number :   60
Comment        :
Number of Volumes :   5
Number of Disks :   1
Names of the Volumes are....
\Device\HarddiskDmVolumes\SQL_DG\REGREP
\Device\HarddiskDmVolumes\SQL_DG\DB_LOG
\Device\HarddiskDmVolumes\SQL_DG\DB_VOL
```

```
\Device\HarddiskDmVolumes\SQL_DG\DATA_FILES  
\Device\HarddiskDmVolumes\SQL_DG\SRL1  
Names of the Disk are....  
Harddisk2
```

The `vx dg` command has a new option, `vx dg refreshff`, which refreshes the state of disk groups on a system to support the new Deported Read-Only state. You must run this command if the disk group state on a passive node does not display the Read-Only state even though FastFailover is enabled for the VMDg resources. This typically occurs in cases where a new node is added to a service group where FastFailover is already enabled.

Enabling fast failover for disk groups in a VCS environment

To support the fast failover feature, a new attribute, `FastFailOver`, is added to the VCS Volume Manager Disk group (VMDg) agent. This attribute decides whether or not a disk group is enabled for fast failover. The `FastFailOver` attribute can take values of 1 and 0. The value 1 indicates that the agent enables fast failover for the configured disk group. The default value 0 indicates that fast failover is disabled for the disk group.

Refer to the *Cluster Server Bundled Agents Reference Guide* for more information about the VMDg agent.

Note: The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click **Properties**. Disk group upgrade is required if you have upgraded SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

To enable the `FastFailover` attribute for a VMDg resource

- 1 In Cluster Manager (Java Console), select a service group with VMDg resource configured for it. Select the **Properties** tab from the right pane.
- 2 Scroll down to choose the **FastFailOver** attribute and click to edit the attribute value.
- 3 In the **Edit Attribute** dialog box, check the **FastFailOver** check box and then click **OK**.
- 4 Repeat these steps for every VMDg resource in the service groups.

Refer to the *Cluster Server Administrator's Guide* for more information about the Cluster Manager (Java Console).

Enabling fast failover for disk groups in a Microsoft Failover Clustering environment

Enabling fast failover for SFW disk groups in a Microsoft Failover Clustering environment involves the following tasks:

- Configuring SFW fast failover support using the SFW Configuration Wizard for Microsoft Failover Cluster
See [“Using the SFW Configuration Wizard for Microsoft Failover Cluster to configure fast failover”](#) on page 173.
- Enabling FastFailover attribute for SFW VMDg resources
See [“To enable fast failover for SFW VMDg resources”](#) on page 174.

To support the fast failover feature, a new attribute, **FastFailOver**, is added to the Volume Manager Disk Group (VMDg) resource. This attribute decides whether or not a disk group is enabled for fast failover. The **FastFailOver** attribute can take values of True or False. The value True indicates that the agent enables fast failover for the configured disk group. The default value False indicates that fast failover is disabled for the disk group.

Note: Apart from configuring SFW fast failover support for Microsoft Failover Clustering, the SFW Configuration Wizard for Microsoft Failover Cluster also prepares a cluster for Cluster Volume Manager (CVM) support. Therefore, it can also be used for configuring CVM in a Microsoft Failover Clustering environment.

Note: The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click **Properties**. Disk group upgrade is required if you have upgraded SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

Enabling fast failover for disk groups in a Microsoft Failover Clustering environment**Using the SFW Configuration Wizard for Microsoft Failover Cluster to configure fast failover**

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), type **Run**, and then press **Enter** to open the **Run** dialog box.

In the **Run** dialog box, type **scc**, and then click **OK** to launch the Solutions Configuration Center (SCC).

- 2 From the Solutions view in the SCC, click on **SFW Configuration Wizard for Microsoft Failover Cluster** to launch the wizard.

- 3 The SFW Configuration Wizard for Microsoft Failover Cluster panel displays cluster and system (node) information and lets you configure or unconfigure systems for fast failover as described below.

Fields	Description
Cluster Name	The automatically-assigned name for the cluster configuration. This name is the same as that of the Microsoft failover cluster.
System Name	This table lists all the systems (nodes) that can participate in the cluster. Select the required cluster nodes that you want to configure for fast failover.
Configuration Status	The following configuration statuses are displayed: <ul style="list-style-type: none"> ■ Configured: The node is configured in the cluster. ■ Not Configured: The node is not configured in the cluster. ■ Invalid Configuration: There are configuration issues with the node. ■ Not Supported: Unsupported configuration. This is displayed if you don't select the Microsoft Failover Cluster option during the product installation.
Configuration Status Details	Displays the configuration status information for the selected systems.

Select one or more systems that are not yet configured and click **Configure** to configure fast failover support for the selected systems.

To unconfigure fast failover support, select one or more configured systems and click **Unconfigure** to unconfigure fast failover support for the selected systems.

- 4 Click **Done** to exit the wizard.

To enable fast failover for SFW VMDg resources

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

Launch the Failover Cluster Manager snap-in by clicking **Failover Cluster Manager**.

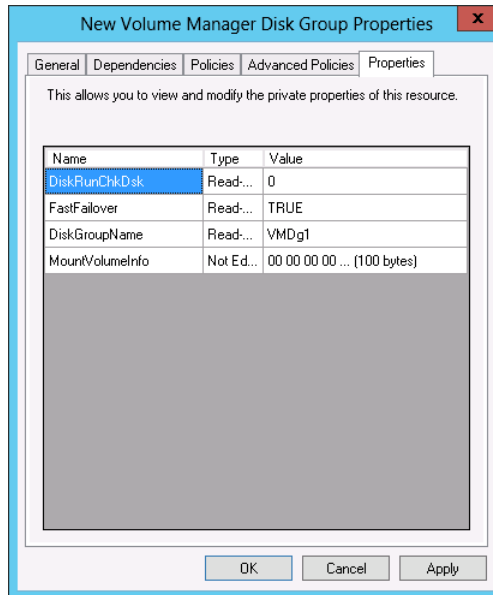
- 2 Right-click **Services and Applications**. Select **More Actions > Create Empty Service or Application**. This creates a service group, for example, **SG1**.

- 3 Right-click the new service group and select **Add a Resource > More Resources > Add a Volume Manager Disk Group** from the context menu.

A new Volume Manager Disk Group (VMDg) resource is created with a default name.

- 4 Right-click the VMDg resource and select **Properties**.

Complete the following on the **Properties** window:



- Select the **General** tab to change the default name of the **New Volume Manager Disk Group** to a name of your choice. Say, for example: **VMDg1**
- Now select the **Properties** tab and perform the following steps:
 - In the **DiskGroupName** box enter the dynamic cluster disk group name created earlier in this document. Say, for example **GuestDG1**.
 - Edit the **FastFailOver** attribute value to **TRUE**.

Displayed default value is **FALSE**.

Note: To enable live migration, you must set the **FastFailOver** attribute to **TRUE** for all VMDg resources in a cluster.

- 5 Right-click the VMDg resource and select **Bring this resource online** option from the center pane of the Failover Cluster Manager snap-in.

Limitations

The following limitations are applicable for fast failover in this release:

- Failover may take longer time to complete when subdisk move operation is in progress
- Failover on iSCSI storage may take long time after NIC pull
If storage and network is configured on the same NIC, then during a NIC pull event, it is observed that the disk groups and the application resources take a longer time to go offline resulting in delayed service group failover.
Recommendation: Configure your storage and network over two independent network adapters. There should be one additional NIC, which is not configured under VCS, connected to storage network.
- Majority Disk Loss scenario in case of partial storage disconnect may take around 30 seconds to declare the disk group as faulted after disk removal events are seen in the VEA console.
- Fault detection takes 20 - 25 seconds in case of active node restart or crash

Configuration and functionality not supported

Fast failover is currently not supported for the following:

- RAID-5 volumes
- SCSI-2
Disk group failovers are not affected. However, the disk groups do not use the fast failover feature if configured using SCSI-2.
- Active/Passive (A/P) arrays for DMP
Only A/PF, A/A , and ALUA arrays are supported.

iSCSI SAN support

This chapter includes the following topics:

- [Overview of iSCSI SAN support](#)
- [Configuring an iSCSI SAN with SFW](#)
- [Managing an iSCSI SAN with SFW](#)

Overview of iSCSI SAN support

An iSCSI SAN is a storage array network that uses the iSCSI protocol that is built on top of TCP/IP for block level I/O. In an iSCSI SAN the basic unit of data transport is called a protocol data unit (PDU). The PDU is transmitted over a session with one or more TCP/IP connections. Each TCP/IP connection is called a portal and is a member of a portal group.

Data stored in an iSCSI SAN is stored on iSCSI targets. An iSCSI target can be a hardware array or a software iSCSI target node such as Microsoft's iSCSI Software Target. An iSCSI initiator is a server that initiates requests to and receives responses (I/O) from an iSCSI target. An iSCSI session is started when an iSCSI initiator sends a logon request to an iSCSI target. The iSCSI target then responds with a logon response. To complete the login, a series of key=value pairs are negotiated and agreed upon. Logins can be made persistent so that the associated storage is available to the iSCSI initiator following a restart.

Challenge Handshake Authentication Protocol (CHAP) and Internet Protocol Security (IPSec) can be used to make the transmission of commands and data secure in an iSCSI SAN. CHAP can be used with IPSec to ensure an authorized connection between an iSCSI initiator and an iSCSI target.

About iSNS servers

An Internet Storage Name Service (iSNS) Server processes iSNS registrations and queries from iSNS clients. It maintains a database to manage the registrations.

The role of the iSNS Server in an iSCSI SAN is to allow iSCSI initiators and iSCSI targets to register and to query for other registered iSCSI initiators or iSCSI targets.

A Discovery Domain (DD) is a security and management mechanism that iSNS Servers use to administer access and connectivity to iSCSI initiators or iSCSI targets (storage nodes) in an iSCSI SAN.

A Discovery Domain Set (DDS) is a mechanism to store sets of discovery domain mappings in an iSNS database. A discovery domain set may contain one or more discovery domains and a discovery domain can be a member of one or more discovery domain sets. Multiple discovery domain sets may be active at the same time.

For the purpose of query and registration, discovery domains are viewed as containers for storage nodes and portal objects. A discovery domain is active if it is a member of at least one active discovery domain set. Discovery domains that are not members of an enabled discovery domain set are considered disabled. A storage node can be a member of one or more discovery domains. An enabled discovery domain establishes connectivity among the storage nodes in that discovery domain.

At the storage node level, a discovery domain is used to make iSCSI targets inaccessible to the iSCSI initiators that are not members of the discovery domain. When first identified by the iSNS server, the Default DD in an iSCSI SAN contains all storage nodes except ones already assigned to a discovery domain.

SFW enables you to register with Microsoft's iSNS service and to manage discovery domains and discovery domain sets.

Configuring an iSCSI SAN with SFW

The Veritas Enterprise Administrator (VEA) GUI provides support for managing an iSCSI SAN.

On an iSCSI initiator node, SFW enables you to define iSCSI target portals, log into and out of iSCSI targets, and view the target portal group configuration. With SFW, you can also connect to and manage iSNS objects on a Microsoft iSNS server. If your server's iSCSI initiator is connected to a Microsoft iSCSI Software Target, you can view the set of LUNs that are accessible to your initiator and assign LUNs on the Target to any of its defined targets from an initiator.

Considering prerequisites

When configuring an iSCSI SAN with SFW, you should consider the prerequisites and the requirements (minimum) for each as described in the documentation that accompanies them:

The following are prerequisites and requirements:

- For Windows Server 2008:
 - VDS 1.1 is included with these versions.
- Microsoft iSCSI initiator 2.06 or above.
- VDS iSCSI Hardware Provider (Microsoft)
(optional; available as part of Windows Storage Server).

Setting up an iSCSI SAN

Setting up an iSCSI SAN requires configuring target portals, setting up the iSCSI targets and iSCSI initiators, configuring the storage, assigning access rights, and registering with an iSNS server.

Setting up the iSCSI targets and configuring the storage should be done according to the instructions of the manufacturer of the storage device.

Setting up the iSCSI initiators requires that iSCSI initiator software be installed on each server to enable them to connect to the iSCSI targets. Optionally, the iSNS server software needs to be installed on a server to allow automatic discovery of the iSCSI targets on the network.

If the iSCSI target that you plan to connect to is not displayed in the VEA GUI under the iSCSI Initiators node, you need to add a target portal for it. This occurs when iSNS is not used or the iSCSI target is not registered with iSNS. This also occurs when there is an iSCSI HBA that is not configured to connect to a particular target.

Assigning access rights

The following shows how to assign access rights.

To assign access rights

- 1 From the tree view of the VEA GUI on the left pane, click the iSCSI node.
If DHCP has been configured to report iSNS servers, the iSNS server automatically discovers the initiators and targets. Expanding the iSCSI node displays all the available initiators and targets in the network. If DHCP has not been configured, you need to register the initiator with iSNS to display the available initiators and targets.
- 2 Logon to the required targets to make them available to the initiator.

Making a target available

The following shows how to make a target available.

To log in to a target and make them available to initiators

- 1 Select a target and select Login from its context menu.
- 2 Check any desired optional login settings. The available login settings are to allow persistent restore of the login or to enable multi-path login.
- 3 To set any security settings, such as CHAP or IPsec, check Advanced Settings to access the security settings dialog.

Make sure that the security settings are compatible with the settings that were set up for the storage device.

Using SFW vxdg lateststart for iSCSI

SFW provides for the dynamic disk groups that contain iSCSI storage to be available at system startup. SFW automatically does this by configuring the Veritas DG Delayed Import Service (VxDgDI) and issuing the `vxdg lateststart` command when a dynamic disk group is created. This automatic configuration of the dynamic disk group is only done on the system where the dynamic disk group is created. SFW also automatically checks the dynamic disk group for persistent login to the iSCSI target where the iSCSI disks are located. If necessary, SFW configures the iSCSI target where the iSCSI disks are located for persistent login.

If this dynamic disk group is moved to another system, you must manually configure the VxDgDI service, issue the `vxdg lateststart` command, and set up iSCSI target persistent login (if required) to make the iSCSI storage available at system startup on the new system.

You can do this by making the VxDgDI service dependent on the iSCSI service (MSiSCSI) in the Windows registry and issuing the `vxdg lateststart` command. You must also change the VxDgDI service startup type to automatic. This allows SFW access to the iSCSI storage when it becomes available. You must also

manually configure the iSCSI target where the iSCSI disks are located for persistent login.

See “[vxdg](#)” on page 593.

In addition to automatically enabling the latestart feature for the dynamic disk groups that are created, SFW automatically maintains or disables the feature as appropriate for the following operations:

- Destroy disk group.
- Remove disk from disk group.
- Add disk to disk group.
- Upgrade dynamic disk group version.
- Split disk group.
- Join disk group.

Managing an iSCSI SAN with SFW

Storage Foundation allows components of the iSCSI SAN to be accessed and managed with the VEA GUI.

Managing with the VEA Console

When an iSCSI SAN has been configured and the Microsoft iSCSI initiator service is installed, SFW displays an iSCSI Initiators node, an iSNS node, and an iSCSI Arrays node in the tree view of the VEA Console.

When managing an iSCSI SAN, you should be already familiar with the usage and settings described in the documentation that accompanies each of the pieces of software that these nodes represent.

About the iSCSI initiators node

Double-clicking the iSCSI Initiators node in the tree-view displays the initiators on the local computer. Selecting one of the nodes enables the display of the initiator adapters, target portals, and the targets for the initiator. These are displayed by clicking the appropriate tab in the right pane. Fully expanding the iSCSI Initiators node in the tree-view displays the target portals and the targets for each initiator in the tree-view.

Selecting a target portal in the tree view displays information about the target portal in the right pane.

Selecting a target node in the tree view displays information about the target portal groups and the sessions of the target in the right pane. Clicking on the portal groups tab displays details about the portal. Clicking on the sessions tab displays details about the connection between the target and initiator. Fully expanding the target node in the tree-view displays the target portals, LUNs associated with the target, and the target sessions for the target.

The following functions to manage the initiators are available in a context menu for each initiator node:

- Configure iSNS Server.
 - Add
 - Remove
- Add Target Portal
- Set Initiator shared secret.
- Set All IPsec Tunnel Mode Addresses.
- Set IPsec Group Pre-Shared Key.
- Properties

The following functions to manage the targets portals are available in a context menu for each target portal:

- Remove Target Portal.
- Properties

The following functions to manage the targets are available in a context menu for each target node:

- Login To Target
- Log out From Target
- Remove Persistent Login
- Properties

The following functions to manage the sessions are available in a context menu for each session node:

- Add Connection
- Logout
- Properties

Details for these functions are as follows:

Configuring an iSNS server	To register the initiator with an iSNS server, right-click the initiator node and select the Configure iSNS Server to open the dialog box. The dialog box lets you add or remove the initiator from an iSNS server and to enter the IP address or DNS name of the server.
Adding a target portal	To add a target portal to an initiator, right-click the initiator node, select Target Portal, and select Add to launch the Add Target Portal Wizard. The wizard also lets you enter the security settings for communicating with the target.
Setting up security for an initiator	<p>You can set up security for an initiator.</p> <p>Use the following to set up the security for an initiator:</p> <ul style="list-style-type: none">■ Set Initiator Shared Secret■ Set All IPsec Tunnel Mode Addresses■ Set IPsec Group Pre-Shared Key <p>To access the appropriate dialog box, right-click the initiator node and select the type of security. To enable persistent login to the target, check Persist in the Set All IPsec Tunnel Mode Addresses dialog box.</p>
Removing a target portal	To remove a target portal from an initiator, use the Remove Target Portal dialog box. To access the dialog box, right-click the appropriate target portal node and select Remove Target Portal. Click Yes to complete the operation.
Logging on to a target and setting security	<p>To log on to a target, right-click the target node, select Login To Target, to launch the Login to Target Wizard. The wizard also lets you enter the login settings and security settings for communicating with the target.</p> <p>The available login settings are to allow persistent login and to enable multi-path login.</p> <p>Check the Advanced Settings checkbox to open the dialog box for security settings.</p> <p>The following security settings are available:</p> <ul style="list-style-type: none">■ General Settings tab:<ul style="list-style-type: none">■ Initiator Adapter■ Target Portal■ CHAP Logon Settings■ Perform mutual authentication■ CRC/Checksum Data digest and/or Header digest■ IPsec Settings tab:<ul style="list-style-type: none">■ Pre-Shared Key■ IKE protocols■ ESP protocols■ Perfect Forward Secrecy

Logout from a target	To logout of a target, use the Logout From Target dialog box. To access the dialog box, right-click the target node and select Logout From Target.
Removing persistent login from a target	<p>Enabling persistent login for a target causes the target to be logged in automatically across system restarts. Removing or disabling persistent login means that manual login to the target is required after a restart.</p> <p>To remove the persistent login from a target, use the Remove Persistent Login dialog box. To access the dialog box, right-click the target node and select Remove Persistent Login.</p>
Adding a connection	<p>To add a connection to a session, right-click the session node and select Add Connection to launch the Add Connection window. The window lets you select initiators and portals to establish the connection.</p> <p>Note: Not all iSCSI initiators, such as the Microsoft iSCSI Initiator, let you add connections.</p>
Logout from a session	To logout of a session, use the Logout dialog box. To access the dialog box, right-click the session node and select Logout

About multiple sessions




Storage Foundation supports multiple sessions to a target. Multiple sessions are used to support fault tolerance for the iSCSI storage using the Microsoft Multipath I/O (Microsoft MPIO) solution, and the Dynamic Multi-Pathing component of the InfoScale products. To set up fault tolerance for the iSCSI storage, install the DMP device-specific modules (DSMs) and VITARGET, and configure the connections to the sessions.

See [“DMP overview”](#) on page 399.

About the iSCSI arrays node

If a VDS iSCSI hardware provider for the storage array is installed on the SFW server, then expanding the iSCSI array node displays the Microsoft iSCSI Software Target subsystem that contains the LUNs and physical disks that make up your iSCSI SAN. The LUNs in a subsystem can be one of the following types shown in the following table.

Table 7-1 Description of LUNs

Icon	Name	Description
	LUN	LUN contained in the subsystem and assigned to this iSCSI initiator.
	Foreign LUN	LUN contained in the subsystem and assigned to another iSCSI initiator.
	Idle LUN	LUN contained in the subsystem, but not assigned to any iSCSI initiator.

Selecting a subsystem displays the tabs that let you view lists of the targets, LUNs, or drives in the right pane.

You can view additional information as follows:

- Double-clicking an entry in the list of targets displays the tabs that let you view the sessions and the portal groups of the target.
- Double-clicking an entry in the list of LUNs displays the tabs that let you view the paths that connect to the physical disk, details of the physical disk, and the targets that the LUN supports.
- Double-clicking an entry in the list of drives displays its details.

Selecting a LUN under the subsystem node displays the tabs that let you view the paths that connect to the physical disk, details of the physical disk where the LUN is located on the target server, and targets.

Expanding a LUN in the tree-view that is assigned to the iSCSI initiator, displays the physical disk for the LUN. Selecting the physical disk displays details of the physical disk in the right pane.

To support an SFW Dynamic Multi-Pathing (DMP DSMs) environment, the MPIO option of the Microsoft iSCSI Initiator and the appropriate DSM for the iSCSI target need to be installed. (SFW provides the VITARGET DSM to support Microsoft iSCSI Target.)

In an SFW Dynamic Multi-Pathing environment, selecting a physical disk under the subsystem node displays the tabs that let you view the volumes, disk regions, disk views, and alerts associated with the physical disk.

The following functions to manage the targets are available in a context menu for each LUN node:

- Assign To Target
- Unassign Target
- Properties

Details for these functions are as follows:

Assign To Target To assign a LUN to a target, right-click the LUN node (under the tree view of the target, under the tree view of the array, or under the LUNs tab when selecting an array), select Assign To Target to launch the Assign LUN To Target Wizard. The wizard lets you select from a list of available targets.

Unassign Target To unassign a LUN from a target, right-click the LUN node (under the tree view of the target, under the tree view of the array, or under the LUNs tab when selecting an array), select Unassign Target to launch the Unassign LUN From Target Wizard. The wizard lets you select from a list of assigned targets.

About the iSNS server node

The iSNS server node appears in the tree view of the VEA GUI on the left pane after performing the Configure iSNS Server option to register the Microsoft iSCSI initiator with the iSNS server you want to manage.

Clicking the iSNS server node displays the discovery domains (DD), discovery domain sets (DDS), and the targets and initiators registered with the server.

Generally speaking, discovery domains, and discovery domain sets are a way to organize and control access to various initiators and targets. Discovery domains contain a collection of initiators and targets. Discovery domain sets are a collection of discovery domains.

Selecting one of the iSNS server nodes enables the display of discovery domains and discovery domain sets. These are displayed by clicking the appropriate tab in the right pane.

Selecting one of the discovery domain nodes enables the display of the set of targets, initiators, and portal groups assigned to the discovery domain. These are displayed by clicking the appropriate tab in the right pane.

Selecting one of the discovery domain set nodes enables the display of discovery domains for the discovery domain set. These are displayed in the right pane.

The following functions to manage the iSNS servers are available in a context menu for each iSNS server node:

- Login
- Log out
- Create Discovery Domain
- Create Discovery Domain Set

The following functions to manage the discovery domains are available in a context menu for each discovery domain node:

- Add node
- Remove node
- Delete

The following functions to manage the discovery domain sets are available in a context menu for each discovery domain set node:

- Add Discovery Domain
- Remove Discovery Domain
- Delete
- Disable
- Enable

Details for these functions are as follows:

Login to an iSNS server	To view or manage an iSNS server you must first login to the iSNS server. To login to an iSNS server, use the Login to iSNS Server dialog box. To access the dialog box, right-click the iSNS server node and select Login.
Log out of an iSNS server	To log out from an iSNS server, use the Log out from iSNS Server dialog box. To access the dialog box, right-click the iSNS server node and select Log out.
Creating discovery domains	To create a discovery domain, use the Create Discovery Domain dialog box. To access the dialog box, right-click the iSNS server node and select Create Discovery Domain. Select from the available iSCSI nodes to add to the discovery domain.
Creating discovery domain sets	To create a discovery domain set, use the Create Discovery Domain dialog box. To access the dialog box, right-click the iSNS server node and select Create Discovery Domain Set. Select from the available discovery domains to add to the discovery domain set.

Adding nodes to a discovery domain	Adding a node to a discovery domain adds an iSCSI initiator or an iSCSI target to the discovery domain. To add nodes to a discovery domain, use the Add Node dialog box. To access the dialog box, right-click the discovery domain node and select Add Node. Select from the available nodes to add to the discovery domain.
Removing nodes from a discovery domain	Removing a node from a discovery domain removes an iSCSI initiator or an iSCSI target from the discovery domain. To remove nodes from a discovery domain, use the Remove Node dialog box. To access the dialog box, right-click the discovery domain node and select Remove Node. Select from the available nodes to remove from the discovery domain.
Deleting a discovery domain	To delete a discovery domain, use the Delete dialog box. To access the dialog box, right-click the discovery domain set node and select Delete.
Adding discovery domains to a discovery domain set	To add discovery domains to a discovery domain set, use the Add Discovery Domain dialog box. To access the dialog box, right-click the discovery domain set node and select Add Discovery Domain. Select from the available discovery domains to add to the discovery domain set.
Removing discovery domains from a discovery domain set	To remove discovery domains from a discovery domain set, use the Remove Domain dialog box. To access the dialog box, right-click the discovery domain set node and select Remove Discovery Domain. Select from the available discovery domains to remove from the discovery domain set.
Deleting a discovery domain set	To delete a discovery domain set, use the Delete dialog box. To access the dialog box, right-click the discovery domain set node and select Delete.
Disabling a discovery domain set	Disabling a discovery domain set makes all the discovery domains in the discovery domain set not available for discovery by other registered targets and initiators in the iSCSI SAN. To disable a discovery domain set, use the Disable dialog box. To access the dialog box, right-click the discovery domain set node and select Disable.
Enabling a discovery domain set	Enabling a discovery domain set makes all the discovery domains in the discovery domain set available to other registered targets and initiators in the iSCSI SAN. To enable a discovery domain set, use the Enable dialog box. To access the dialog box, right-click the discovery domain set node and select Enable.

Settings for monitoring objects

This chapter includes the following topics:

- [Overview](#)
- [Event monitoring and notification](#)
- [Event notification](#)
- [Disk monitoring](#)
- [Capacity monitoring](#)
- [Configuring Automatic volume growth](#)
- [SMTP configuration for email notification](#)

Overview

Storage Foundation (SFW) provides different ways to monitor objects.

Event monitoring and notification

Storage Foundation (SFW) reports the storage subsystem events that you can view in the Alert Log. The log can provide information on significant incidents, such as a disk failure. SFW also provides event notification by SMTP email, by pager, and through SNMP traps that can be displayed in HP OpenView, IBM Tivoli, and CA Unicenter.

See [“Alert log”](#) on page 189.

See [“Severity levels”](#) on page 189.

See “[Log configuration](#)” on page 189.

See “[Log configuration](#)” on page 189.

Alert log

Click the Logs folder in the Logs Perspective to display the Alert Log, a listing of alerts that pertain to the Storage Foundation for Windows program. For each alert listing, you see information about the date and time of the message, the message text, and its class.

The most recent alert listings can also appear in the new section at the bottom of the Veritas Enterprise Administrator (VEA) GUI. This happens when the Console tab is selected from the bottom left corner of the VEA GUI. If the Tasks tab is selected, text appears in the bottom section of the window that gives the progress of system tasks, such as format, refresh, and rescan. These items also appear in the Logs Perspective under the Task Log tab.

You can view specific information on each event by double-clicking on it to display the Alert Log Details dialog box. This shows an alert description, recommended action, and properties. The dialog box provides an easier way to read the description of the event. The arrows that appear at the upper-right corner can be used for scrolling in the alerts listing.

Severity levels

The alerts that are listed in the Alert Log have severity levels.

The severity levels are as follows:

- Critical — Indicates alerts of high importance involving a system crash or unrecoverable data loss.
- Error — Indicates a loss of functionality or data.
- Warning — Indicates recoverable errors and no data corruption.
- Information — Indicates an informational event, such as a new disk has been added to the system.

Log configuration

You can configure the Alert Log and Task Log through the Log Settings dialog box. To access this dialog box, select the Control Panel Perspective, select the host you are connected to, and double-click the **Log Settings** icon in the right-pane.

In this dialog box, you can enter the maximum log size for each of the logs. You can access the Event Log through the VEA console.

Event notification

Storage Foundation for Windows provides event notification by SMTP email, by pager, and through SNMP traps that can be displayed in HP OpenView, CA Unicenter, and IBM Tivoli. You can configure the notification service to send messages to specific individuals or groups and also set up rules for filtering the event messages. You can also set up rules that run a script when specified conditions apply.

This section presents an overview of the Rule Manager utility, which provides the event notification service for Storage Foundation for Windows. The full details for this utility are given in the online Help module titled "Optional Rules for Handling Events." To access the online Help, select Contents from the SFW Help menu.

See ["Accessing the Rule Manager utility"](#) on page 190.

See ["Preliminary setup"](#) on page 190.

See ["Steps for creating rules"](#) on page 191.

See ["Accessing the SFW SNMP MIB files"](#) on page 193.

Accessing the Rule Manager utility

The commands for configuring the service are all accessed through the Rule Manager icon. If you click Control Panel in the Perspective bar and expand the tree view, the actionagent node is displayed. Clicking the actionagent node displays the Rule Manager icon in the right pane.

Double-click the **Rule Manager** icon to bring up the **Rule Manager** window. All functions relating to event notification rules are initiated from this window.

From the Rule Manager, you can create, enable, disable, modify, and delete rules relating to event notification. The New Rule button at the top of the window is used to create new rules. The Configure Mail button is used to configure settings, such as the settings for the Mail Server, Default Sender, and SMTP Port.

Preliminary setup

Before creating rules, you must configure the SMTP mail notification settings you want to use.

Note: Note that this is a one-time setting. After providing this information, you don't need to repeat this step unless the Mail Server changes for your system.

To configure the SMTP mail notification settings

- 1 From the VEA Control Panel perspective, select the **actionagent** node in the tree view.
- 2 Double-click **Rule Manager** in the right pane.
- 3 Click **Configure Mail**.
The Mail Settings window is displayed.
- 4 Provide information about your Mail Server, Default Sender, and SMTP Port. The default SMTP Port is 25.
Check the Enable check box and click **OK**.

Steps for creating rules

After the preliminary setup, you create rules for notifications.

To create rules

- 1 From the VEA Control Panel perspective, select the **actionagent** node in the tree view.
- 2 Double-click **Rule Manager** in the right pane.
The Rule Manager window is displayed.
- 3 Click **New Rule** to start the New Rule Wizard.
The New Rule Wizard Welcome appears.
- 4 Review the steps of the wizard and click **Next** to continue.
- 5 On the **Rule Properties** window, do the following:
 - Enter a name for the rule.
 - Select the type of rule you want to create. Choose one of the following:
 - A rule based on the Alert topic.
 - A rule based on the severity of alerts.
 - Further define your rule by selecting either an Alert topic or a severity level as follows:
 - Enter a topic or click **Select** to select a topic from a list if you chose a rule based on an Alert topic.
 - Select one of the severity levels if you chose a rule based on the severity of the alerts.Click **Next** to continue.

- 6 On the **Select host for Rule** window, select **Trigger rule for alerts originating from specific hosts** and check the appropriate host in the lower part of the window.

Click **Next** to continue.

- 7 On the **Select Actions** window configure one or more of the following actions to be taken when the events are detected:

- Send email notification to recipients.
Enter the email addresses that are to receive notification.
- Send SNMP trap notification to a host.
Specify the host name and port number to use to send the notification.
- Execute command on a host.
Specify the name of the host, source host that has the source file (script), path of the source file, and the parameters for the source file. The source file can be of any executable file type, such as `.cmd` or `.bat`; and the source file must not be dependent on any libraries. Separate the parameters by spaces or by the delimiter that the source file recognizes. In standalone mode, the host and source host identify the same host.

An individual log file is created and maintained for each rule that executes a command on a host. Each time the command is executed entries are made into its log file. The log files are located at `C:\Documents and Settings\All Users\Application Data\Veritas\VRTSaa\Log`.

Click **Next** to continue.

- 8 On the **Finish Rule Setup** window, review your settings and check the **Enable Rule** check box.

Click **Finish** to complete setting up the rule.

About SNMP traps

If you plan to send SNMP traps, the MIB files should be loaded and compiled into the management application to which traps are sent.

To accomplish this, do the following:

- Load the product-specific MIB into your SNMP management software database according to the instructions provided by your management platform. See [“Accessing the SFW SNMP MIB files”](#) on page 193.
- If the platform you’re using requires a configuration file, such as `trapd.conf`, you should load that file as well. Consult the user guide for the management platform to determine how to map your traps to the platform’s alarm system.

Any configuration files that are associated with your product are stored in the same location as the MIB files.

Note: For details on creating rules, see the "Optional Rules for Handling Events" online Help module. To access this module, select Contents from the Help menu.

Accessing the SFW SNMP MIB files

If you want to receive Storage Foundation for Windows SNMP traps on a management platform (such as HP OpenView), you need to compile and load the Veritas SNMP MIB files into your management platform, following the procedures outlined in the documentation for your specific platform.

When you install Veritas InfoScale Foundation, Veritas InfoScale Storage or Veritas InfoScale Enterprise, the Veritas MIB files are installed on a local hard drive along with the Veritas Object Bus files and the SFW installation files. (The Veritas Object Bus is associated with the VEA console and is needed to run Storage Foundation for Windows.)

The default path to the MIB files directory is: C:\Program Files\Veritas\VERITAS Object Bus\Providers\actionprovider\win-snmpmibfiles.

The Storage Foundation for Windows MIB file for the Veritas Object Bus is VRTS-vm.mib and can be found in the snmpmibfiles directory.

Disk monitoring

Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.) is available on many new SCSI disks. S.M.A.R.T. provides a way for disks to self-check certain physical conditions that are useful predictors of disk failure and send an alert when conditions indicate that a failure may occur soon. This allows the system administrator to take preventive action to protect the data on the disk. To receive alerts in the Storage Foundation for Windows Event Log, S.M.A.R.T. monitoring must be enabled for the global configuration and also on each disk you want checked. The default is to have S.M.A.R.T. monitoring disabled.

If S.M.A.R.T. predicts that a failure occurs, an alert is sent to the Event Log. You can also receive event notification by SMTP email, by pager, and through the SNMP traps that can be displayed in HP OpenView, IBM Tivoli, and CA Unicenter.

See ["Event notification"](#) on page 190.


Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support S.M.A.R.T. monitoring.

To enable S.M.A.R.T. monitoring

- 1 Right-click on a disk for which you want to enable S.M.A.R.T. monitoring.
- 2 Select **S.M.A.R.T. Monitoring** from the context menu.

If the **S.M.A.R.T. Monitoring** command on the context menu is grayed out, then the S.M.A.R.T. monitoring technology is not available on that disk.

A S.M.A.R.T. monitoring dialog box appears as shown below.



- 3 If it is not already checked, check the **Enable S.M.A.R.T. monitoring** check box in the global configuration.

The **Enable S.M.A.R.T. monitoring** check box acts as an on/off switch for the entire system and lets you temporarily disable S.M.A.R.T. monitoring if desired.

- 4 Select the rate in hours for S.M.A.R.T. monitoring to occur.

Since this is a global configuration, the monitoring rate is the same for all disks in the system. The default rate is one hour.

S.M.A.R.T. monitoring uses system resources as it checks the device.

- 5 Check the **Enable S.M.A.R.T. check on this device** check box at the bottom of the window to have S.M.A.R.T. monitoring be performed on the specific disk you selected.
- 6 Click **OK**.
- 7 Repeat steps 1, 2, 5, and 6 for each disk you want to monitor.

Capacity monitoring

Capacity monitoring is a feature that monitors the capacity of Storage Foundation for Windows dynamic volumes and Microsoft Disk Management dynamic volumes so that when the used disk space on a volume reaches a user-specified threshold, an email alert message is sent.

You can also use SFW's Rule Manager event notification function to set up an email message that is sent when thresholds are reached and logged.

See [“Accessing the Rule Manager utility”](#) on page 190.

Capacity monitoring is set at the individual volume level and is turned off by default. You need to enable it for each volume you want to monitor.

The capacity monitoring settings you establish only affect existing volumes. New volumes that are added to a server after any capacity monitoring settings are made have the default capacity monitoring settings applied to them. If the default is not appropriate, you must change the settings for any new volumes that are added.

Capacity monitoring is also available in the CLI with `vxvol capacitymonitor`.

To enable a volume warning notification

- 1 Right-click on the volume you want to manage and click **Capacity Monitoring**. The Capacity Monitoring window appears.

localhost - Capacity Monitoring Wizard

Enable capacity monitoring for volume FSVOL1 (F:)

Disable capacity monitoring for volume FSVOL1 (F:)

Polling interval for checking capacity (sec):

Minimum time between alerts if condition persists (sec):

Capacity Monitoring Thresholds:

Critical threshold (%):

Warning threshold (%):

Send email notification when volume has reached its growth or warning threshold

Recipient E-mail Address(es):

Make policy available to another cluster node when failover

2 Set the capacity monitoring parameters.

To set the capacity monitoring parameters, perform the following:

Enable Capacity Monitoring for Volume	Check this option to enable capacity monitoring for a selected volume
Disable Capacity Monitoring for Volume	Check this option to disable capacity monitoring for a selected volume
Polling interval for checking capacity (sec)	The polling interval is the interval of time, in seconds, at which Storage Foundation checks volume capacities. The minimum and default value is 20 seconds.
Minimum time between alerts	To keep the event log file from getting full, you can limit how often an alert message is generated after the warning threshold is reached for a volume. The default is 100 seconds between messages.
Critical threshold	The percentage of volume capacity that Storage Foundation reports an error. The default critical threshold is 90% of capacity.
Warning threshold	The percentage of volume capacity that Storage Foundation logs a warning alert message. This is a warning that the volume is reaching full capacity. The default warning threshold is 80% of capacity.
Send email notification when volume.....	Enable the check box marked against " Send email notification when volume..... " if you want to send email notifications. In the Recepient Email Address (es) , enter the email IDs to whom notification needs to be sent.
Make Policy available to another cluster node when failover	Enable the check box Make Policy available to another cluster node when failover if you have a cluster setup. This check box appears only in case of a clustered setup.

3 Click **OK** to enable the settings.

Configuring Automatic volume growth

The Automatic Volume Growth feature is a command to have an individual volume grow (increase capacity) automatically by a user specified amount when the used

disk space on the volume reaches a certain threshold. This threshold is set in the Capacity Monitoring wizard.

By automating the growth process, the Automatic Volume Growth feature enhances the capability of dynamic volumes to expand without requiring that the underlying disks be offline.

The advantages of this automatic process include the following:

- Reduced need for a System Administrator to monitor volume capacity at all times.
- Disk space is allocated on an as-needed basis; no idle excess capacity.
- Space is always available for mission-critical data.

Automatic Volume Growth is set at the individual volume level and is turned off by default. You need to enable it for each volume you want to grow. Enabling a volume for Automatic Volume Growth requires the volume to be enabled for Capacity Monitoring.

The following gives more information about configuring Automatic Volume Growth in a cluster environment.

See [“Enabling Automatic Volume Growth in a non-clustered and clustered environment”](#) on page 199.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Automatic Volume Growth.

A policy file is used to record the settings for the volumes that use the Automatic Volume Growth feature. This XML file is automatically set up when Automatic Volume Growth is enabled, and it reflects the current settings for a volume. You can view and change the current settings with the Automatic Volume Growth wizard.

Alternately, you can use SFW’s Rule Manager event notification function to grow a volume automatically. With Rule Manager, you can set up a script that can grow a volume automatically when a specified threshold is reached.

The current Automatic Volume Growth settings of all your volumes are also available to help you manage your volumes. Selecting the volume node in the tree-view of the GUI and selecting Automatic Volume Growth Information in the context menu, displays the growth size, the maximum growth limit, and if Automatic Volume Growth is enabled for your volumes.

The following topics give more information on Rule Manager:

See [“Automatic volume growth script for volume capacity monitoring”](#) on page 201.

See [“Accessing the Rule Manager utility”](#) on page 190.

Automatic Volume Growth is also available in the CLI and with `vxvol autogrow`.

Warning: If you use a script to enable volume growth, do not enable the Automatic Volume Growth function through the Automatic Volume Growth wizard. If you enable a volume growth script and enable the Automatic Volume Growth function for a volume, the volume increases in size by the combined amount. If you prefer to use the new feature, you must disable the script.

Enabling Automatic Volume Growth in a non-clustered and clustered environment

Automatic volume growth feature can be enabled on a non-clustered and clustered environment as shown in the below procedure.

Note: In a cluster environment, you must configure the Veritas Scheduler Service on all nodes to allow the transfer of Capacity Monitoring and Automatic Volume Growth settings to other nodes after failover. These settings are done once for each node after the product installation or upgrade.

See [“Configuring the Veritas Scheduler Service on a clustered environment”](#) on page 200.

To enable the Automatic Volume Growth feature, perform the following steps:

To enable automatic volume growth in a non-clustered and clustered environment

- 1 Right-click on the volume you want to manage and click **Automatic Volume Growth**.
The Automatic Volume Growth wizard appears.
Click **Next** to continue.
- 2 Highlight any additional volumes you want to manage and click **Add**.
Click **Next** to continue.
- 3 Check the **Enable automatic volume growth** check box.
- 4 Set the Automatic Volume Growth parameters as follows:
 - Volume growth
Enter the size by which you want the volume to expand each time the growth threshold is reached.
 - Maximum volume size
Select the maximum volume size for the volume:

Unrestricted volume growth	Default setting. Allows the volume to grow without an upper limit to the volume size.
Restricted volume growth	Lets you enter a maximum size for the volume. The volume growth is limited to this specified size. This setting may be useful to manage a database volume that may be vulnerable to uncontrolled rapid growth due to malicious updates, spam, or virus attack.

- 5 Enable notification by email (optional) as follows:
 - **Send email notification when volume has reached its growth threshold**
Check this check box, if you want an email sent when a volume has reached its growth threshold.
Enter the email address(es) to where the notifications should be sent in the Recipient E-mail Address(es) box. Separate multiple addresses with commas.
Make sure that the SMTP Configuration settings have been set. If they have not, you should set them now.
See “[SMTP configuration for email notification](#)” on page 204.
- 6 Enable the check box **Make policy available to another cluster node when failover**, if you have a cluster environment and want the automatic volume growth policy to be available to another cluster node after a failover.

Note: This check box is available only when the volume belongs to a cluster disk group.

- 7 Click **Next** to continue.
- 8 Select which disks to automatically grow the volume as follows:
 - **Auto select disks**
SFW automatically selects which disks to grow.
 - **Manually select disks**
Highlight which disks you want to grow and click **Add**.
Click **Next** to continue.
- 9 Review the summary of the Automatic Volume Growth settings and click **Finish** to apply.

Configuring the Veritas Scheduler Service on a clustered environment

To configure the Veritas Scheduler Service on cluster nodes, perform the following steps:

To configure the Veritas Scheduler Service in a cluster environment

- 1 Create a domain user that has administrator privileges that can access every cluster node.

Note: For Windows Server operating systems, the user must be the domain administrator.

Note: For Exchange 2010, the user account that is configured for the Scheduler Service should have read-write permissions to the file share path in case of a VCS cluster setup.

- 2 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.
Under **Services**, select **Veritas Scheduler Service**.
- 3 Right-click **Veritas Scheduler Service** and select **Properties** from the context menu.
- 4 Click the **Log On tab** on the **Properties** window.
- 5 Click the **This Account option** and enter the domain user/administrator ID and password that was created in an earlier step of this procedure.
Confirm the password and click **Apply**, then **OK**.
- 6 On the Windows Services GUI, restart the Veritas Scheduler Service to complete the configuration.

Automatic volume growth script for volume capacity monitoring

The volume capacity monitoring feature can send alert messages when dynamic volumes reach certain size thresholds. In addition, with the use of a script, volume capacity monitoring can automatically grow a volume when it reaches a specified capacity. With a script, you can conserve disk space on your servers because space is distributed automatically on an as-needed basis. You do not have to manually allocate the additional disk space when it is required.

The following steps set up a rule for growing a volume when a Capacity Monitoring threshold error occurs. The volume capacity monitoring feature monitors the capacity of the volume on your system; it has two thresholds—a Warning threshold with a default setting of 80% of capacity and a Critical threshold with a default setting of 90% of capacity. These are the points at which the volume capacity monitoring feature sends an alert. To automate volume growth, you can associate an executable

file, `volgrow.cmd`, with a Rule Manager rule so that when the Critical threshold is passed, a command to expand the volume is automatically sent.

Before setting up a new rule to execute a command, you must enable the VEA service and the GridNode service to interact with the desktop. These steps are not required if the new rule sends only email notifications.

Enabling the VEA and GridNode to interact with the desktop

The following describes how to enable the VEA service and the GridNode service with the desktop.

To enable the VEA service and the GridNode service to interact with the desktop

- 1 Stop the following services in the order shown below before configuring a new rule:
 - Action Agent
 - Storage Agent
 - Veritas Enterprise Administrator
 - GridNode
- 2 Right-click the Veritas Enterprise Administrator service. Select Properties, click **Allow service to interact with desktop** under the Log On tab, and click **OK**.
- 3 Right-click the GridNode service. Select Properties, click **Allow service to interact with desktop** under the Log On tab, and click **OK**.
- 4 Start the following services in the order shown below after configuring a new rule:
 - GridNode
 - Veritas Enterprise Administrator
 - Storage Agent
 - Action Agent

Configuring a rule to execute volume growth

The following describes how to configure a rule to execute volume growth when capacity thresholds are met.

To configure a new rule to execute volume growth when capacity thresholds are met

- 1 Make sure that the volume capacity monitoring feature is turned on for the volume you want to monitor and that the thresholds are set to the desired percentages. The volume capacity monitoring feature is off by default.
The following enables volume capacity monitoring:
 - Right-click the volume and select **Capacity Monitoring**.
 - Select enable capacity monitoring for the volume.
 - Change the appropriate settings.
You may want to set the **Minimum time between alerts if condition persists** to a large number, such as 3600 seconds (equivalent to one hour), so that you won't get so many messages while the volume is enlarged.
 - Click **OK** after the settings are completed.
- 2 If you have not already done so, set up the recipient or recipients that receive the alert messages about the capacity threshold.
See ["Event notification"](#) on page 190.
- 3 Select the Control Panel perspective in the VEA GUI.
- 4 Click the **actionagent** node in the tree view.
- 5 Double-click the **Rule Manager** icon to display the Rule Manager window.
- 6 Click the **New Rule** button to display the Rule wizard.
Click **Next** to continue.
- 7 In the window that appears, do the following:
 - Enter a name for the rule and select **A rule based on the alert topic**.
 - Click **Select** to select the Alert topic.
 - Check the **Volume capacity reached error condition** alert check box to select the alert.
 - Click **OK** to close the windows and return to the wizard.Click **Next** to continue the wizard.
- 8 Specify if the rule is triggered for alerts from any host or from specific hosts. If you specify alerts originating from specific hosts, then check the appropriate hosts in the lower part of the window.
Click **Next** to continue.

- 9 Check both check boxes in the **Select actions** window to enable the following actions:

■ **Send Email Notification**

Enter the email addresses of the recipients.

■ **Execute Command**

Enter the name of the hosts, the source host, and the path of the source file.

Enter the path for the `volgrow` command file for the source file. The default path for `volgrow.cmd` is:

```
C:\Program Files\Veritas\Veritas Volume Manager\  
Sig\volgrow.cmd <VolumeName|DriveLetter>
```

The included `volgrow.cmd` script is set to grow a volume by 150 MB. The size for growing the volume can be changed by modifying the `volgrow.cmd` script to a desired number. Please read the comments in the script for more information.

Click **Next** to continue.

- 10 Check **Enable Rule** and click **Finish** to complete the setup.

The Rule Manager window displays listing the new rule. With the new rule, once the amount of data in the volume reaches the point where the error threshold is reached, the rule activates, a message is sent that the error threshold has passed, and the command to grow the volume is implemented immediately. Once the process for growing the volume is complete, you stop receiving error messages about the volume capacity.

SMTP configuration for email notification

The SMTP Configuration settings are global settings and need to be set only once for all the capacity monitoring and automatic volume growth email messages that are sent.

To configure SMTP

- 1 From the VEA Control Panel Perspective, select the StorageAgent node in the tree view.
- 2 Select **SMTP Configuration** in the right pane. The SMTP Settings screen appears.
- 3 Enter the following information:

- SMTP server: Enter the name or IP address of the SMTP mail server on your network.
 - Port: The default port for the SMTP mail server is usually 25.
 - Sender: Enter the name that should appear as the sender.
 - Sender email address: Enter the full SMTP email address for the sender.
- 4** Click **OK** to save your settings.

Standard features for adding fault tolerance

This chapter includes the following topics:

- [Fault tolerance overview](#)
- [Mirroring](#)
- [RAID-5 and RAID 0+1](#)
- [Hot relocation](#)
- [Dirty region logging \(DRL\) and RAID-5 logging](#)
- [Dynamic relayout](#)

Fault tolerance overview

Fault tolerance is the capability of ensuring data integrity when hardware and software failures occur. This chapter focuses on the Storage Foundation for Windows standard features that provide fault tolerance.

Note: Additional fault tolerance for your storage can be gained through the program's options. The Microsoft Failover Clustering support option provides fault tolerance for up to eight servers through clustering. The Dynamic Multi-Pathing option provides multiple paths from the server to a storage array to ensure application availability.

Mirroring

This section describes how to work with mirrored volumes.

See [“About mirroring”](#) on page 207.

See [“Add a mirror to a volume”](#) on page 208.

See [“Two methods for creating a mirror”](#) on page 208.

See [“Break or remove a mirror ”](#) on page 210.

See [“Set the mirrored volume read policy”](#) on page 212.

See [“Mirror a dynamic boot or system volume”](#) on page 214.

See [“Reactivate a mirrored volume”](#) on page 214.

See [“Repair a mirrored volume”](#) on page 215.

For advanced usages of mirrored volumes:

See [“Dirty region logging for mirrored volumes”](#) on page 222.

See [“FastResync”](#) on page 253.

In addition, the FastResync feature is an advanced procedure for mirrored volumes. FastResync is a feature that is part of the FlashSnap option.

Note: FlashSnap is available only with the Enterprise edition. See the Licensing section in the *Veritas InfoScale Installation and Upgrade Guide* for details.

See [“FastResync”](#) on page 253.

About mirroring

A mirrored volume consists of at least two subdisks of identical sizes that are located on separate disks. Each time a file system or application writes to the volume, Storage Foundation for Windows transparently writes the same data to each of the volume's subdisks. With Storage Foundation for Windows, you can have up to 32 mirrors for all volume types except for RAID-5.

Add Mirror: The primary purpose of mirroring is to provide fault tolerance; however, mirrored volumes can also improve I/O performance for most I/O-intensive applications, which make substantially more read requests than writes. With mirrored volumes, write requests may take longer than with non-mirrored volumes, but the effects are typically minor.

Break Mirror: An important use of mirroring is the practice of breaking off a third mirror that can be used as a snapshot for backup or for other activities, such as uploading or updating data warehouses or performing application testing.

Remove Mirror: Removing a mirror from a volume "removes" or destroys the data from the selected mirror and leaves the other mirror or mirrors intact. After you

remove a mirror, the disk space that the removed mirror uses becomes unallocated free space.

You use the Mirror Wizard to complete these tasks. The end of the Mirror Wizard displays a summary of the details of the operation.

Two methods for creating a mirror

You can create a mirror on an existing volume or create a new mirrored volume.

For details, see the following:

- To create a mirror on an existing volume, use the Add Mirror command. See [“Add a mirror to a volume”](#) on page 208.
- To create a new mirrored volume, use the New Volume wizard. You select a mirrored layout when you create the volume. See [“About creating dynamic volumes”](#) on page 47.

Note: The mirroring process is disabled if the volume cannot be mirrored or if there is no unallocated space on another dynamic disk large enough to mirror the volume. A RAID-5 volume cannot be mirrored.

Add a mirror to a volume

The basic process of adding a mirror is described in this section.

The Add Mirror wizard also enables you to implement the Dynamic Relayout feature. See [“Dynamic relayout”](#) on page 225.

Note: If InfoScale Foundation, InfoScale Storage, or InfoScale Enterprise is installed on Windows Server 2008 R2, then the Add Mirror operation for dynamic volume gets disabled if this operation is performed through the Logical Disk Management (LDM) console. In such cases, it is recommended to perform the Add Mirror operation through the Veritas Enterprise Administrator (VEA) console.

If you create more than one mirror at a time, you may see inconsistent information on the progress bar. Also, the generation of multiple mirrors does affect system resources. After creating a mirror, you may want to wait until it has finished generating before creating another mirror.

Warning: Adding mirrors to a volume involves a certain amount of time for mirror resynchronization. Take care not to perform actions on the volume until the command is complete. Also, if you inadvertently shut down the server or deport the disk group containing the volume before mirror resynchronization completes, any mirrors that are in the process of being added to the volume are deleted when the disk group with the volume comes online again.

Note that the SmartMove feature can help reduce mirror resynchronization time.

See “[SmartMove](#)” on page 93.

To add a mirror to a volume

1 Right-click on the volume you want to mirror.

2 Select **Mirror>Add** from the context menu.

The Add Mirror wizard appears.

3 Specify how many mirrors you want to add to the volume, and, if desired, manually assign the destination disks for the mirrors.

To assign the disks manually, click the **Manually select disks** radio button. Then use the buttons for adding and removing disks to move the disks you want assigned as destination disks to the right pane of the window.

You may also check **Disable Track Alignment** to disable track alignment for the mirrors.

Click **Next** to continue.

4 In the **Specify the attributes** dialog box, select the Concatenated or Striped volume layout type. If you create a striped mirror, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.

You may also specify a mirror that mirrors across disks with the following attributes:

- Connected by port
- Identified by target
- Contained in an enclosure
- Connected by channel

In addition for a striped layout, you may specify to have the stripe across disks done by the following:

- Port
- Target

- Enclosure
- Channel

Click **Next** to continue.

The operation to create a new dynamic volume fails if the appropriate resources are not available to support the selected attributes to mirror across disks.

- 5 On the summary page, click **Finish** to add the mirror.

Break or remove a mirror

You can break or remove a mirror. It is important to understand the difference between these operations.

Breaking a mirror takes away a redundant mirror (or plex) of a volume and assigns it another drive letter. The data on the new volume is a snapshot of the original volume at the time of breaking. Breaking off a plex of the mirrored volume does not delete the information, but it does mean that the plex that is broken off no longer mirrors information from the other plex or plexes in the mirrored volume.

Removing a mirror from a volume "removes" or destroys the data from the selected mirror and leaves the other mirror or mirrors intact. After you remove a mirror, the space on the disk used by the removed mirror becomes unallocated free space.

Note: Storage Foundation for Windows does not let you remove the mirror of the volume that Windows uses to boot from. Also, if you break a mirrored boot volume, you cannot boot from the portion that was broken off.

The details on breaking and removing a mirror are included in the next two sections, "Break Mirror" and "Remove Mirror."

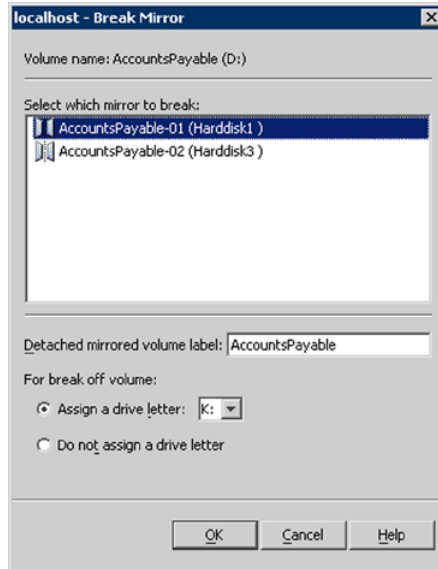
Break mirror

This section explains how to break a mirror.

Note: If you have a mirrored volume with two plexes and you break one off, the remaining plex is no longer a mirrored volume. It too assumes its other layout characteristics without the mirroring, but it does keep the drive letter of the original mirrored volume.

To break a mirror

- 1 Right-click on the volume from which you want to break a mirror. Click Mirror on the context menu that comes up, and then select Break from the submenu.



- 2 In the **Break Mirror** dialog box that appears, select the mirror you want to break off.

Choose whether or not to assign a drive letter to the broken-off volume. You may assign a specific letter from the drop-down list or accept the default. You can also assign a label to the detached mirror.

When you are satisfied with your choices, click **OK**.

Once the command is completed, the broken-off plex is no longer part of the mirrored volume and is assigned the drive letter specified, if **Assign a drive letter** was selected. The broken-off plex retains all other volume layout characteristics except the mirroring. For example, if you had a mirrored striped volume, the broken-off plex would become a striped volume.

Remove mirror

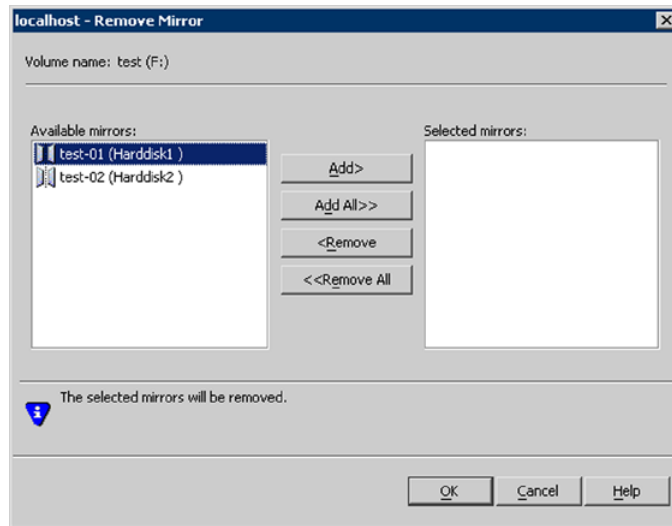
After a mirror is removed, the space formerly allocated to it reverts to free space.

Warning: When a mirror is removed, all of the data on the removed mirror is deleted.

To remove a mirror

- 1 Right-click on the volume from which you want to remove the mirror.
- 2 Select **Mirror > Remove**.

The Remove Mirror dialog box appears as shown below:



The name of the volume with the mirror appears in the upper left.

- 3 The available mirrors appear in the left pane. Use the **Add** or **Add All** button to move the desired mirrors to the list of mirrors selected to be removed in the right pane. To move mirrors from the **Selected mirrors** list to the **Available mirrors** list, select them in the right pane and click the **Remove** button. To move all of the mirrors from the **Selected mirrors** list back to the **Available mirrors** list, click the **Remove All** button.
- 4 Select the mirror to remove.
Click **OK** to remove the mirror or mirrors.
- 5 Click **Yes** at the message warning you that removing mirrors reduce the redundancy of the volume.

Set the mirrored volume read policy

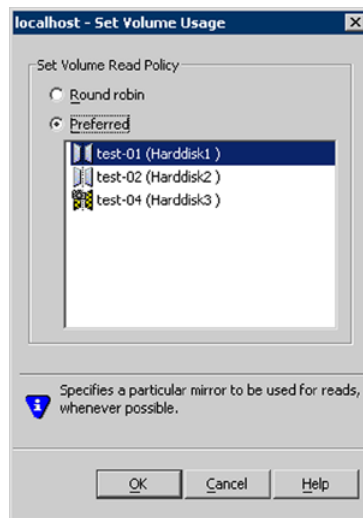
Setting the read policy on a dynamic volume lets you specify either that a particular mirror be used for reads or that all mirrors be read in turn in "round-robin" fashion for each nonsequential I/O detected. The default is that mirrors be read "round robin."

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support setting the volume read policy.

To set the volume read policy

- 1 Right-click on the volume you want to set the read policy for.
- 2 Select **Set Volume Usage** on the context menu.

The Set Volume Usage dialog appears.



- 3 Select either the **Round robin** or **Preferred** option.

These options are as follows:

- The Round robin option 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 takes advantage of the drive or controller read-ahead caching policies.
- If you select the Preferred option, you must select the mirror you want to set as the preferred option from the list of mirrors shown. This mirror is used for reads whenever possible.

This can improve your system's read performance, particularly if the drive used for the target mirror is a high-performance drive, or if the target mirror is located on a drive that is geographically closer.

If a mirror has already been set as preferred, you see the text "(Preferred plex)" after the mirror's name in the dialog box.

When a preferred plex has an error, another plex completes the read operation. This happens automatically, as if there were no preferred plex.

- 4 Click **OK** to set the volume usage, or click the **Cancel** button to exit without setting the read policy.

Mirror a dynamic boot or system volume

This section presents items to consider when you plan to mirror a dynamic boot or system volume.

You need to be aware of the following:

- Be sure to back up any data on the system before you create the mirror.
- For information on how you can boot your system from the mirror for Windows Server operating systems, refer to the Microsoft documentation.
- If a break-mirror operation is performed on a mirrored boot volume, the resulting new volume—the broken-off mirror—is not usable as a boot volume.

Under certain conditions, you may be unable to create a mirror on a boot or system volume or get an error message during the process.

See [“An attempt to mirror a boot or system volume fails or has an error message”](#) on page 549.

You must follow specific steps in setting up a dynamic boot or system volume.

See [“Set up a dynamic boot and system volume \(Optional\)”](#) on page 59.

Reactivate a mirrored volume

If there is a problem with one of the disks on which a mirrored volume resides, make sure that the disk is properly connected, and then try to bring back that disk online by first using Rescan. If Rescan does not work, then use the Reactivate Disk command. In most cases, Reactivate Disk brings the disk online and all of the volumes are healthy.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Reactivate Disk command.

However, if any of the mirrors of the mirrored volume are still not healthy, you may need to resynchronize the degraded volume with the Reactivate Volume command. Resynchronizing makes the data consistent on all mirrors of a mirrored volume.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Reactivate Volume command.

Right-click the volume and select **Reactivate Volume** from the context menu.

The mirrored volume should come back to a Healthy state unless there is a serious problem with the volume.

See also the following topics:

See [“Rescan command”](#) on page 536.

See [“Reactivate disk command”](#) on page 540.

See [“Reactivate volume command”](#) on page 541.

Repair a mirrored volume

When a disk on which a mirrored volume resides fails, the volume displays a Degraded status. The disk's name is changed to Missing Disk, and an icon (X) appears on the Missing Disk icon. The status of the disk is Offline.

You can use the Repair Volume command to repair a mirrored volume.

See [“Repair a mirrored volume”](#) on page 215.

RAID-5 and RAID 0+1

This section describes the fault tolerance capabilities of RAID-5 and RAID 0+1.

RAID-5

A RAID-5 volume is a fault-tolerant volume with data and parity striped intermittently across three or more physical disks. Parity is a calculated value that is used to reconstruct data after a failure. If a portion of a physical disk fails, the data on the failed portion can be recreated from the remaining data and parity. You can create RAID-5 volumes only on dynamic disks. You cannot mirror RAID-5 volumes.

The advantage of RAID-5 is that it provides fault tolerance with less hardware than that required for mirroring. RAID-5 supports the configuration of any number of data disks with the hardware cost of only one disk for parity. Mirroring requires separate disks for each mirror. RAID-5 protects against data loss when one disk fails.

However, it is necessary to resynchronize the data from the other disks to access the data. With mirroring, if one disk fails, the data is available from another mirror right away. Also, RAID-5 does not have the advantage of adding more mirrors for increased fault tolerance. With mirroring, you can have three or more mirrors so that if more than one disk fails, there is still fault tolerance. With RAID-5, once two

disks fail, the RAID-5 volumes on multiple disks fail. Now that the cost of storage has decreased, mirroring is a better choice for fault tolerance.

You can resynchronize a RAID-5 volume after it is degraded (having lost one disk).

See [“Repair volume command for dynamic RAID-5 volumes”](#) on page 542.

RAID 0+1

Mirroring with striping, RAID 0+1, offers the protection of mirroring and the speed of striping.

The following are some advantages of RAID 0+1 volumes:

- Very large fault-tolerant volumes can be created that can hold large databases or server applications.
- Read performance is very high. A preferred plex can be designated. Within each plex, striping balances the load.
- The write penalty (that is, the multiple writes SFW must perform to keep all plexes' contents synchronized) is mitigated by the striping of data across multiple disks.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support RAID 0+1.

About the only disadvantage to RAID 0+1 is the hardware cost. The user must purchase, house, power, and operate twice as much raw storage capacity as the data requires. Today with lower storage costs, mirroring with striping is a preferred RAID layout for mission-critical data.

Hot relocation

'Hot relocation' describes the ability of a system to automatically react to I/O failures. It is an automatic process that relocates subdisks from a disk that has I/O errors.

The default for Storage Foundation for Windows is to have automatic hot relocation mode inactive. This means that if an I/O error occurs in a redundant subdisk, the subdisk is not automatically relocated to another disk. The tasks to enable the hot relocation mode are described in the following hot relocation setup section.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support hot relocation.

Note: Dynamic disks belonging to a SAN Disk Group do not support hot relocation.

See [“Hot relocation overview”](#) on page 217.

See [“Hot relocation setup summary”](#) on page 218.

See [“Hot relocation mode”](#) on page 218.

See [“Customize target disks with set disk usage”](#) on page 219.

See [“Undo hot relocation”](#) on page 220.

See [“Clear hot relocation information”](#) on page 221.

Hot relocation overview

In hot relocation mode, if there are I/O errors on a subdisk containing a redundant volume (RAID-5 or mirrored), the redundant volume’s subdisks are moved from the failed subdisk to a hot spare disk. If there is not enough free space on designated hot spare disks, then free space on any available disk is used. Depending on available disk space, the subdisks is not always moved to the same disk, but may be scattered within the dynamic disk group. The Undo Hot Relocation command can be used to move the subdisks back to their original location after the failed disk has been repaired or replaced.

Note: The hot relocation operation is performed only for redundant subdisks with I/O errors if there are healthy plexes of the volume available elsewhere within the dynamic disk group. Nonredundant subdisks with I/O errors are not relocated; the volume fails, and an alert message is sent to the Event Log.

When selecting space for relocation, hot relocation preserves the redundancy characteristics of the dynamic volume that the relocated subdisk belongs to. For example, hot relocation ensures that a failed subdisk is not relocated to a disk containing its mirror. If redundancy cannot be preserved by using any available spare disks 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.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

See [“Administering disks for site-based allocation”](#) on page 454.

When hot relocation takes place, the failed subdisk is removed from the configuration database, and Storage Foundation for Windows takes precautions to ensure that the disk space used by the failed subdisk is not recycled as free space.

The default is to have all disks available as targets for hot-relocated subdisks. Using the Set Disk Usage command, you can choose to designate specific disks as not available.

See [“Disk usage options”](#) on page 219.

Note: The Reset button restores default settings.

Hot relocation setup summary

There are two tasks that must be done to ensure that the hot relocation mode is properly configured.

The first task uses the Hot Relocation command from System Settings to verify that the hot relocation mode is enabled.

The second task uses the Set Disk Usage command to define within each disk group which disks, if any, should be selected as preferred targets for hot-relocated subdisks or never used as targets for hot-relocated subdisks.

Hot relocation mode

The default for Storage Foundation for Windows is to have automatic hot relocation mode inactive. This means that if an I/O error occurs in a redundant subdisk, the subdisk is not automatically relocated to another disk. Use the Set Disk Usage command to designate preferred disks as targets for hot relocation.

The option to enable the hot relocation mode is available from System Settings in the Control Panel or from Edit Configuration in the Storage Foundation Assistant.

To enable the hot relocation mode

- 1 Open **System Settings** in the Control Panel or open **Edit Configuration** in the Storage Foundation Assistant.
- 2 Click the **Hot Relocation** tab in **System Settings** or click **Hot Relocation** in **Edit Configuration**.
- 3 Check the check box to enable the hot relocation mode. If automatic hot relocation is disabled and an I/O failure occurs, an alert is sent to the Event Log, but no further action takes place.
- 4 When you are satisfied with your selection, click **OK**.

Customize target disks with set disk usage

The Set Disk Usage command lets you specify how Storage Foundation for Windows uses individual disks within a dynamic disk group. If an I/O failure occurs, SFW looks for free space on preferred or secondary target disks within the dynamic disk group that contains the failing disk to determine where to move the subdisks. Thus, it is important to consider the entire disk group when configuring disk usage. In addition, you need to consider the track alignment settings of the failing disk and the target disks. SFW does not preserve the track alignment settings of the failing disk. SFW adopts the track alignment settings of the target disk after hot relocation takes place.

The default for each disk is "secondary hot relocation target." This is the state you prefer for the majority of disks within each disk group. If hot relocation occurs and subdisks from a specific disk are scattered throughout the disk group, it is possible to repair or replace the failed disk and then use the Undo Hot Relocation command to return all the subdisks back to their original location.

Disk usage options

The Set Disk Usage command is used to set disk usage options.

It provides the following options:

- **Reserved for manual use**

Disks reserved for manual use are not available in automatic selection operations, including hot relocation. In an automatic selection operation, Storage Foundation for Windows chooses the storage where the operation occurs. Generally, the user is given a choice between allowing SFW to "Auto select disks" or "Manually select disks." Examples of commands that allow automatic selection are New Volume and Add Mirror.

Reserving a disk for manual use lets you prevent any unwanted volumes or subdisks from being placed on that disk and gives you complete control over the disk.

There are three options related to how a disk is used as a target for hot-relocated subdisks. In considering how to mark each disk, remember that the hot relocation operation occurs at the disk group level, and only target disks within the same disk group are considered.

The three options are as follows:

- Preferred hot relocation target
If there is an I/O failure anywhere in the system, SFW first looks for space on the disks that have been marked as preferred hot-relocation targets for redundant subdisks.
- Secondary hot relocation target
This is the default for all disks. During the hot relocation operation, if there are no disks selected as preferred targets or if there is no space available on those disks, SFW chooses space on disks marked as secondary targets.
- Not used as a hot relocation target
This option does not allow any hot-relocated subdisks to be moved to the selected disks. It differs from the "Reserved for manual use" option in that the disk remains available for other automatic selection operations.

To set disk usage

- 1 Right-click on the desired disk.
- 2 Select **Set Disk Usage** from the context menu.
The **Set Disk Usage** dialog box appears.
- 3 Choose how you want Storage Foundation for Windows to use the disk.
- 4 When you are satisfied with your selection, click **OK**.

You can review the settings for all your disks by selecting the Disks node in the tree view and then viewing the **Disks** tab in the right pane. The settings for each disk are listed in the Reserved and Hot Use columns.

Undo hot relocation

The command Undo Hot Relocation relocates subdisks back to their repaired original disk or replacement disk and restores a system to its original configuration, less any failed volumes. If hot relocation scattered subdisks from a failed disk to several disks within a dynamic disk group, the Undo Hot Relocation command moves all of them back to a single disk without requiring the user to find and move each one individually.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Undo Hot Relocation command.

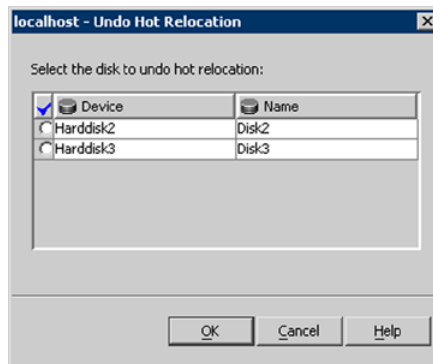
Note: Dynamic disks belonging to a SAN Disk Group do not support the Undo Hot Relocation command.

Note: This command is available only after a hot relocation procedure has occurred.

To undo a hot relocation

- 1 Select the dynamic disk group that contains the hot-relocated volumes.
- 2 Select **Undo Hot Relocation** from the context menu.

A window appears showing the original disks that had I/O errors.



- 3 Click in the column under the checkmark to select the disks you want to restore. The hot-relocated subdisks return to this disk.
- 4 Click **OK**.

Clear hot relocation information

The Clear Hot Relocation Information command deletes the information that describes the original location of any subdisks that have been moved by the hot relocation process. Use this command only if you do not plan to move the hot-relocated subdisks back to their original configuration.

Warning: Once you clear the hot relocation information, you cannot perform an Undo Hot Relocation command for that dynamic disk group.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Clear Hot Relocation Information command.

Note: Dynamic disks belonging to a SAN Disk Group do not support the Clear Hot Relocation Information command.

To clear hot relocation information

- 1 Right-click on a disk group that contains hot-relocated subdisks.
- 2 Select **Clear Hot Relocation Information** from the context menu.
- 3 A dialog box appears asking if you are sure that you want to clear the relocation information for the dynamic disk group. Click **Yes** if you want to clear the hot relocation information.

Dirty region logging (DRL) and RAID-5 logging

Dirty region logging (DRL) and RAID-5 logging greatly speed up the time that it takes to recover from a system crash for mirrored volumes and RAID-5 volumes.

This section describes the dirty region logging (DRL) and RAID-5 logging features.

See [“Dirty region logging for mirrored volumes”](#) on page 222.

See [“RAID-5 logging”](#) on page 223.

See [“Add a log”](#) on page 223.

See [“Remove a log”](#) on page 224.

Dirty region logging for mirrored volumes

DRL uses a log-based recovery method to quickly resynchronize all the copies of a mirrored volume when a system is restarted following a crash. If DRL is not used and a system failure occurs, all mirrors of the volume must be restored to a consistent state by copying the full contents of the volume between its mirrors. This process can be lengthy and I/O intensive.

A log can be created when the volume is created or can be added later.

Warning: Although DRL makes sure that the mirror copies of the mirrored volumes are in sync following a system crash, it does not guarantee data integrity. The file system preserves the data integrity, for example, NTFS.

When a dirty region log is created, a log subdisk is added to the volume. Multiple logs can be associated with a single mirrored volume as a fault-tolerant measure. Depending on the number of logs configured, this may result in degraded performance.

RAID-5 logging

The primary purpose of RAID-5 logging is to quickly reactivate a RAID-5 volume when a system is restarted following a crash. Without RAID-5 logging enabled, a time-consuming reactivation of the entire RAID-5 volume is necessary.

When a RAID-5 log is created, a log subdisk is added to the volume. Multiple RAID-5 logs can be associated with a single RAID-5 volume as a fault-tolerant measure. Depending on the number of logs configured, this may result in degraded performance. If more than one log is added, then the logs are just mirrored. DRL logs for RAID volumes are called RAID-5 logs and provide the same functionality.

If a RAID-5 log is available, updates need to be made only to the data and parity portions of the volume that were in transit during a system crash. The entire volume does not have to be resynchronized.

A log can be created when the volume is created or can be added later.

Add a log

This operation adds a log plex containing a log subdisk to a volume. For mirrored volumes, a dirty region log is added. For RAID-5 volumes, a RAID-5 log is added.

Note: Logs are used only for mirrored volumes or RAID-5 volumes.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support adding a log.

To add a log to a mirrored or RAID-5 volume

- 1 Right-click on the volume you want to add the log to.
- 2 Select **Log>Add** from the context menu.
The Add Log dialog box appears.
- 3 Choose either the default setting **Auto assign destination disks** or the option to manually assign the destination disks. Click **OK** to create the log.

Auto assign destination disks

This option adds one log to the selected volume.

Manually assign destination disks

You can create multiple logs with this option.

Click the **Manually assign destination disks** radio button.

Available disks appear in the left pane. To add destination disks for the logs, select them in the left pane and click the **Add** button. To add logs to all available disks, click the **Add All** button.

To remove destination disks, select them in the right pane and click the **Remove** button. To remove all of the destination disks, click the **Remove All** button.

Disable Track Alignment

This option disables track alignment for the logs being added.

Remove a log

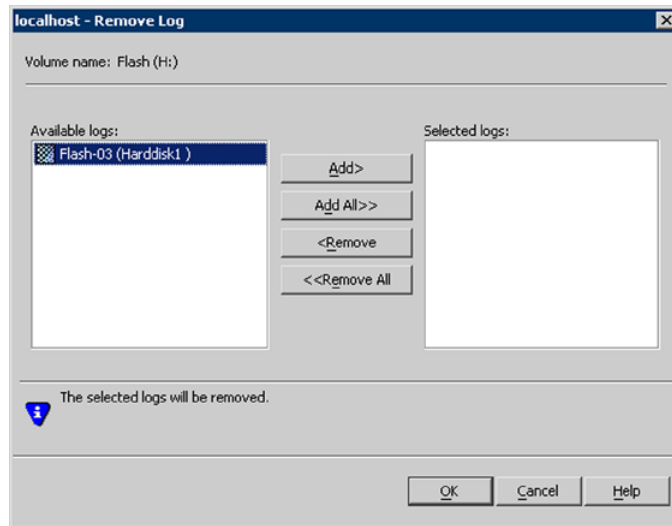
After a log is removed, the space formerly allocated to it reverts to free space.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support removing a log.

To remove a log

- 1 Right-click on the volume you want to remove the log from.
- 2 Select **Log > Remove**.

The Remove Log dialog box appears as shown in the sample screen below:



The name of the volume with the log appears in the upper left.

- 3 The available logs appear in the left pane. Use the **Add** or **Add All** button to move the desired logs to the list of logs selected to be removed in the right pane. To move the logs from the **Selected logs** list to the **Available logs** list, select them in the right pane and click the **Remove** button. To move all the logs from the **Selected logs** list back to the **Available logs** list, click the **Remove All** button.
- 4 After selecting the log, click **OK** to remove the log.
- 5 Click **Yes** at the message warning you that removing logs potentially increase recovery time.

Dynamic relayout

Storage Foundation lets you reconfigure the layout of a volume without taking it offline with the Dynamic Relayout feature.

You can make the following dynamic volume layout changes:

Table 9-1 Dynamic Relayout Changes

Changed from	Changed to
Concatenated	Striped
Striped	Concatenated
Striped	Striped with changes in column size or stripe unit size

This feature is implemented through the Add Mirror window. That window has a section called "Choose the layout."

The basic concept is that you select the features you want the volume layout to be changed to, and the program creates a mirror or plex of the original volume with the desired attributes of the new layout. If you no longer have a need for the original volume, you remove it with the Remove Mirror command.

Thus, the conversion or relayout is a two-step process. First, you select the desired attributes and create the mirror through the Add Mirror command. Second, you remove the original plex through the Remove Mirror command.

Examples of dynamic relayout

This section provides examples of the different layout types. The section that follows gives the step-by-step process for Dynamic Relayout.

The following are the different layout types:

- **Concatenated to Striped**
For example, if you want to convert a concatenated volume on a single disk to a striped volume with two columns, you select striped and two columns in the **Add Mirror** window. The resulting volume has two mirrors or plexes—one is the original plex, and the second is a copy of the original volume striped across two disks. Thus, the entire volume requires a total of three disks. The volume's type is designated as Mixed. At this point, you can remove the original plex by deleting it with the Remove Mirror command. Now you have a striped volume in place of the concatenated one that existed before.
- **Striped to Concatenated**
Changing a striped volume to a concatenated volume probably won't be done very often. Assume the example of an existing two-column striped volume. If you open the Add Mirror window and select the concatenated layout, the resulting volume involves three disks: The first two disks have the plex of the two-column striped original volume, and the third disk contains a concatenated plex of the volume. By removing the original plex, the volume now resides on one disk.
- **Striped to Striped**

The purpose of converting a striped volume to another striped volume is to change the attributes of the volume (the stripe unit size and the number of columns) without having to recreate the volume.

For example, you can convert a two-column striped volume to a three-column striped volume. First, you would select three columns in the **Add Mirror** window. The resulting volume would involve five disks: two disks for the original two-column striped plex and three disks for the new three-column striped mirror. Again, you can remove the original plex by deleting it with the Remove Mirror command, leaving only a three-column striped volume.

Detailed steps for dynamic relayout

Dynamic Relayout lets you change dynamic volume layout types while the volumes are still online. You use both the Add Mirror command and the Remove Mirror command to complete the Dynamic Relayout process.

To complete the Dynamic Relayout process:

- 1** Select the volume with the layout that you want to converted. Choose **Mirror > Add** from the context menu.
- 2** Set the number of mirrors setting to 1. If desired, you can manually select the disks.

To select the disks and use the **Add** or **Add All** button to move the disks you want to the **Selected disks** pane on the right side of the window.

You may also check **Disable Track Alignment** to disable track alignment for the mirrors.

Click **Next** to continue.

- 3** Select the desired layout for the volume in the "Choose the layout" section of the window.

If you choose striped, make sure that you have the number of columns desired and the stripe unit size that you want.

Click **Next** to continue.

- 4 When you are satisfied with your settings, click **Finish**.

The program creates a mirror or plex of the original volume with the attributes of the new layout. If you no longer have a need for the original volume, you can remove it with **Remove Mirror**.

You may want to review the information on how the program implements the dynamic relayout for each of the possible layout types.

See [“Examples of dynamic relayout”](#) on page 226.

- 5 If you want to delete the original plex, right-click on it and select **Mirror>Remove** from the context menu. In the **Remove Mirror** window that appears, use the **Add** button to move the plex to be removed to the **Selected mirrors** pane on the right side of the window, and click **OK**.

The process for converting the layout type on the volume is now complete.

Performance tuning

This chapter includes the following topics:

- [Overview](#)
- [Statistics overview](#)
- [Real-time statistics collection and display](#)
- [Historical statistics collection and graphing](#)
- [Subdisk move, split, and join](#)
- [Command line commands for statistics](#)
- [SmartMove](#)

Overview

Storage Foundation for Windows (SFW) provides performance tuning for storage objects on dynamic disks.

Statistics overview

The statistics feature of Storage Foundation for Windows provides I/O statistics to let you do performance tuning to improve overall disk and system performance.

Storage Foundation for Windows provides the following statistics parameters:

- **Read Requests/Second**
The number of read requests per second for selected storage objects.
- **Write Requests/Second**
The number of write requests per second for selected storage objects.
- **Read Blocks/Second**

The amount of read request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the read requests made.

- **Write Blocks/Second**
The amount of write request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the write requests made.
- **Average Time/Read Block**
The average time in microseconds that it takes to process a read block from the time a request is made until the data is returned.
- **Average Time/Write Block**
The average time in microseconds that it takes to process a write block from the time a request is made until the data is returned.
- **Read and Write Requests/Second**
The number of read and write requests per second for selected storage objects.
- **Read and Write Blocks/Second**
The number of read and write blocks per second for selected storage objects.
- **Queue Depth**
The current number of read and write requests in the queue for selected disks. It does not apply to volumes and subdisks.

With these statistical tools, you can improve disk and system performance by the following actions:

- Identifying high I/O areas that are known as "hot spots."
- Moving data among physical drives to evenly balance the I/O load among the disk drives. This is known as "load balancing."

Bottlenecks occur when a program or device uses too much disk I/O, creating a hot spot and degrading performance. By monitoring a system's I/O, you can find the problem areas and prevent bottlenecks, thus ensuring smooth operation. The Performance Monitoring utility of Storage Foundation for Windows (SFW) lets you continuously monitor and tune system performance. You can make adjustments by moving subdisks from an area of high I/O usage to another disk that has lower usage. You also have the feature of splitting subdisks and joining them back together again. This gives you added flexibility of breaking subdisks down to smaller parts and then moving them.

You can view the statistics in a real-time format in the **Online Monitoring** window. You can also set up the statistics collection so that you can do an analysis of the statistics over time with the historical statistics collection and graphing functions.

The table below summarizes the similarities and differences between the two types of statistics.

Table 10-1 Comparison of real-time and historical statistics

	Real-time	Historical
Displays in	The Online Monitoring window that shows real-time statistical values in a tabular format for selected storage objects.	A graph that represents historical statistics for one storage object. You can select multiple storage objects to track, but each object is graphed individually.
Statistical Parameters	You can select one or more of nine statistical parameters to display in the Online Data Display Options window.	All nine statistical parameters are included in the data collection, but you select one parameter at a time for graphing from the graphing window.
Storage Objects	When you select Online Monitoring, the Select Objects for Online Monitoring window comes up first to let you select the storage objects to monitor.	You select storage objects to monitor in the Start Historical Data Collection window.
Time Frame	Current values since last refresh. Data is not saved. Online monitoring data is available only when the Online Monitoring window is open.	Data is accumulated over time and saved in a file. You indicate the file size. You can start and stop data collection. You view the data only through a graph.

For details on these two types of statistics formats, see the following topics:

See [“Real-time statistics collection and display”](#) on page 231.

See [“Historical statistics collection and graphing”](#) on page 237.

Real-time statistics collection and display

This section describes real-time statistics collection and its display, which is done in the **Online Monitoring** window.

See [“Online monitoring window features”](#) on page 231.

See [“Setup steps for real-time statistics”](#) on page 233.

Online monitoring window features

The Online Monitoring window displays real-time statistics for storage objects you select. It can include disks, subdisks, and volumes. To access the Online Monitoring

window, select Statistics>Online Monitoring from the Tools menu, then select the storage objects to monitor in a preliminary window that comes up, and click **OK**.

The Online Monitoring window collects statistics as long as you have it open. The window can operate in the background. You can do other commands while the window is open. When you set up the real-time statistics, you can change the refresh rate, the time interval of the refresh of the screen display. The default is 5 seconds. The statistics are computed continuously in the background.

The Online Monitoring window is organized as follows:

- **Name**
 The first column contains the names of the storage objects that have been selected for monitoring. Thus, each row contains the statistical data on a selected storage object.
- **Hot Spot**
 The second column has a pie symbol that is designed to be used as an indicator for hot spot detection. A hot spot is an area of high I/O activity that may cause bottlenecks in I/O throughput. The color and shading of each pie symbol indicate the current state of I/O activity for the associated storage object. The pie symbol for the highest amount of I/O activity, Critical activity, is red and fully shaded, making it easy to notice on the screen. You can select the statistical parameters you want to use for hot spot detection.
- **Statistical Parameters**
 The remaining columns represent the different statistical parameters available in the program. By default, Read Requests per Second, Write Requests per Second, Read and Write Requests per Second, Read and Write Blocks per Second, and Queue Depth are selected, but you can select any of the nine statistical parameters.
 See “[Statistics overview](#)” on page 229.
- **Legend for the Hot Spot Indicator**
 The legend at the bottom of the window describes the four pie symbol categories for the hot spot indicator. The table below gives more detail on these categories. If the pie symbol indicator displays the High or Critical category, it signals the fact that you have a hot spot—that is, an area of high I/O activity.

Table 10-2 Explanation of the pie symbol categories for hot spot indicator

I/O Activity	Pie Symbol	Subdisk I/O	Dynamic Disk I/O
Low	Blue Small section in pie	Does not apply	More than 30% below average if it is not below minimum I/O threshold

Table 10-2 Explanation of the pie symbol categories for hot spot indicator
(continued)

I/O Activity	Pie Symbol	Subdisk I/O	Dynamic Disk I/O
Normal	Green Section less than half of pie	80% or less of parent disk I/O	30% below average up to 30% above average
High	Yellow Section more than half of pie	81% to 90% of parent disk I/O	31% to 50% above average
Critical	Red Color in all of pie	More than 90% of parent disk I/O	More than 50% above average

- Context Menu

If you right-click any row in the **Online Monitoring** window, a context menu appears with the following two menu choices relating to the storage object being monitored:

- Graph Historical Statistics Information
- Remove Object from Online Monitoring

The first menu choice is grayed out unless you have previously selected that storage object for historical statistics monitoring.

Setup steps for real-time statistics

This section describes what is involved in setting up the real-time statistics, so that you can view them in the **Online Monitoring** window.

See [“Getting acquainted with the statistics submenu”](#) on page 233.

See [“Selecting online display options”](#) on page 234.

See [“Selecting storage objects to monitor”](#) on page 236.

Getting acquainted with the statistics submenu

The commands for setting up the real-time statistics are located on the Statistics submenu on the Tools menu.

The table that follows summarizes the available commands:

Table 10-3 Table of statistics commands

Command	Description
Online Data Display Options	Sets display options for real-time statistics.
Online Monitoring	Brings up the Select Objects to Monitor window to let you select objects to monitor for a real-time statistics session. Once you have selected objects to monitor, the Online Monitoring window appears.
Start Historical Data Collection	Starts the data collection for historical statistics.
Stop Historical Data Collection	Stops the data collection for historical statistics.

Historical statistics are displayed as graphs of individual storage objects that you select.

See [“Historical statistics collection and graphing”](#) on page 237.

Selecting online display options

The first task in the setup process is to set the display options. This task is not necessary. If you do not use this window, the default settings, which are shown in this window the first time you bring it up, apply. The default settings are shown in the sample screen below.

To select the online data display options

- 1 Choose **Statistics>Online Data Display Options** from the **Tools** menu to bring up the Online Data Display Options window.
- 2 (Optional) Change the **I/O Refresh Rate/Sec** setting.
This is the time interval of the refresh of the screen display. The program continuously runs the actual statistics in the background but changes the display of them in the **Online Monitoring** window according to the refresh rate you set. The default and minimum I/O refresh rate is 5 seconds.
- 3 (Optional) Change the **Threshold** setting.
The threshold is the minimum amount of I/O per second that a disk or subdisk has to have to display a color other than blue. Blue is the lowest state, low I/O activity. The default threshold is 15.
- 4 Determine which statistical parameter or parameters to use for the **hot spot indicator**.

The window provides the following two places to enter a hot spot indicator:

- The drop-down list that is labeled "Hot spot indicator based on."
- "The check box labeled "Use queue depth as hot spot indicator."

Queue depth applies only to disks. The other eight statistical parameters that are in the "Hot spot indicator based on" drop-down list apply to disks, volumes, and subdisks.

If you select the check box to use queue depth as a hot spot indicator, queue depth is used to calculate the hot spot statistic for disks only. The indicator that is used to calculate the hot spot statistic for volumes and subdisks is the statistical parameter that is shown in the "Hot spot indicator based on" drop-down list.

If you don't select queue depth as a hot spot indicator, then the indicator that is used to calculate the hot spot statistic for volumes, subdisks, and disks is the statistical parameter that is shown in the "Hot spot indicator based on" drop-down list.

Why Use Queue Depth?

It is recommended that you select queue depth when you have both disks that are directly connected to your computer and disks that are in a hardware disk array. Because a disk in a disk array can have a much higher level of reads or writes and still function well, reads or writes per second are not really a good way to compare directly connected disks and array disks. Queue depth is the current number of read and write requests in the queue for selected disks. It is a better indicator of how well a disk in an array works, and it also works for the disks that are directly attached to the computer.

- 5 Select the statistics parameters you want to monitor by clicking the check box in front of each parameter. The statistics parameters are described as follows:
 - Read Requests/Second
The number of read requests per second for selected storage objects.
 - Write Requests/Second
The number of write requests per second for selected storage objects.
 - Read Blocks/Second
The amount of read request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the read requests made.
 - Write Blocks/Second
The amount of write request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the write requests made.
 - Average Time/Read Block

The average time in microseconds that it takes to process a read block from the time a request is made until the data is returned.

- **Average Time/Write Block**
The average time in microseconds that it takes to process a write block from the time a request is made until the data is returned.
- **Read and Write Requests/Second**
The number of read and write requests per second for selected storage objects.
- **Read and Write Blocks/Second**The number of read and write blocks per second for selected storage objects.
- **Queue Depth**
The current number of read and write requests in the queue for selected disks. It does not apply to volumes and subdisks.

By default, Read Requests per Second, Write Requests per Second, Read and Write Requests per Second, Read and Write Blocks per Second, and Queue Depth are selected parameters. The **Reset** button restores default settings.

- 6 When you are finished with your changes, click **OK**.

Selecting storage objects to monitor

When you select the Online Monitoring command, the **Select Objects for Online Monitoring** window appears to let you select the storage objects to monitor. The detailed steps are given in the procedure that follows.

To select objects for online monitoring

- 1 Choose **Statistics > Online Monitoring** from the **Tools** menu, or click the **Online Monitoring** button from the Storage Foundation Assistant.

The **Select Objects for Online Monitoring** window appears. You can select an entire disk group or its individual disks, subdisks, and volumes.

You may need to click on the plus symbol next to a disk group icon to display the disks, subdisks, and volumes.

- 2 Select the storage objects you want to display by using the following options:
 - **Add**: Adds one or more selected objects in the left pane of the window to the **Selected List** in the right pane of the window. If you select a dynamic disk group and click **Add**, all associated objects (disks, subdisks, and volumes) are added to the **Selected List** in the right pane.
 - **Add All**: Adds all disks, subdisks, and volumes to the **Selected List** in the right pane.

- **Remove:** Removes any objects that are selected from the **Selected List** in the right pane.
 - **Remove All:** Removes all available objects from the **Selected List** in the right pane.
- 3 Click **OK** to view the statistics in the **Online Monitoring** window.
See [“Online monitoring window features”](#) on page 231.

Historical statistics collection and graphing

The historical statistics collection and graphing function lets you monitor one or more storage objects (disks, subdisks, or volumes) over time and to create individual graphs for these storage objects. You invoke the graphing function by right-clicking a storage object and then selecting the command Graph Historical Statistics Information from the context menu that appears. However, the command is not available on the context menu without certain setup steps having been done first.

This section first describes the graphing window and then details the setup functions that are necessary before you can do the graphing.

See [“Features of the graphing window”](#) on page 237.

See [“Setup steps for historical statistics”](#) on page 238.

See [“Stopping historical data collection”](#) on page 241.

Features of the graphing window

After you select the Graph Historical Statistics Information command from a storage object’s context menu, a graphing window comes up, but it does not display a graph because no statistical parameter has been selected. To display the statistics, you select the View menu at the top left corner of the graph and click to choose one of the nine statistical parameters to graph.

The features of the graph include the following:

- The vertical axis displays a scale for the statistical parameter being graphed, and the name of the statistical parameter is also displayed vertically next to the axis.
- The horizontal axis shows time, which the data collection rate setting determines. The name of the storage object being graphed is displayed under the axis.
- The View menu has the various statistical parameters, letting you change the graph to show a different parameter.

- The File menu has an Exit command, but you can also close the graph window by clicking the close box in the top right corner.

A graph can be saved and printed by copying it to the Clipboard with Alt+PrintScreen and then pasting it into a program, such as Microsoft Word.

Setup steps for historical statistics

This section describes how to configure and graph historical statistics.

See [“Configuring settings in the historical statistics settings window”](#) on page 238.

See [“Selecting storage objects and starting the data collection for historical statistics”](#) on page 239.

See [“Graphing steps”](#) on page 240.

Configuring settings in the historical statistics settings window

The first task in the setup process is to configure the settings in the **Historical Statistics Settings** window. This window is accessed by clicking StorageAgent in the tree view and then clicking Control Panel in the Perspective bar. If you accept the default settings, which are the settings that are appropriate for most data collection purposes, then you do not need to use the following procedure.

To access the historical statistics settings window to view and adjust settings

- 1** In the VEA console tree view, click the **StorageAgent** node and then click **Control Panel** in the Perspective bar.

The Control Panel icons appear in the right pane of the window, including the **Historical Statistics** icon.
- 2** Double-click the **Historical Statistics** icon to bring up the Historical Statistics Settings window.
- 3** Check the settings in the **Historical Statistics Settings** window (as described below), and make changes if desired. Click **OK** to confirm the settings and close the window.

Note that the **Reset** button restores default settings.

Historical Statistics Window Settings

This section describes the settings for historical statistics.

The are as follows:

- Data collection rate in seconds

This is the rate at which the data is collected. It is the time interval between data samples. The default is every 5 seconds. You may prefer the data to be collected every hour (3,600 seconds). This setting establishes the point when the data is graphed. If you plan to run the statistics for several days, it is easier to read the graph with a longer data collection rate.

- **File size per object monitored**
Storage Foundation for Windows collects the data in a file. You set the desired size of the file. Once the file reaches the specified size, the older statistics information is deleted as new information is added. This feature lets you continuously monitor the statistics without taking up increasing file space.
- **Disk Usage Information section**
The numbers at the bottom of this window are an estimate of how long (in days and hours) it would take to compile statistics per storage object monitored to fill up the file size indicated. If you change either the data collection rate or the file size per object monitored, click the Check button to update the estimate.

Selecting storage objects and starting the data collection for historical statistics

The next task in the setup for the historical statistics is to select the storage objects you want to monitor and start the historical statistics data collection. These functions are available through the Start Historical Statistics command. This task is not optional.

To select storage objects and start the historical data collection

- 1 Select **Tools > Statistics > Start Historical Data Collection** to bring up the Start Historical Data Collection window, as shown in the sample screen.

This window lets you select the storage objects for historical data collection. Once you have made your selection and clicked **OK**, the historical data collection begins. The data collection continues in the background until you stop it with the **Stop Historical Data Collection** command or until Storage Foundation for Windows is stopped or the computer is restarted.

- 2 Select the storage object or objects you want to monitor for historical statistics data collection by highlighting the object or objects and using the **Add** button to move the objects to the **Selected List**.

You can also move the objects to the right pane by double-clicking them or by dragging them to the right pane. Use the **Remove** or **Remove All** button to move items on the **Selected List** in the right pane back to the **Devices** list in the left pane.

- 3 Click **OK** when you have made your selection.

The data collection begins in the background. You now have the settings necessary for invoking the graphing function. If you select a storage object that was one of the selected objects or that is included under a selected object (for example, if you selected an entire disk group), the **Graph Historical Statistics Information** command will be available on that storage object's context menu.

Graphing steps

This section covers the specific steps for creating a graph.

You must have selected one or more storage objects and have started the historical data collection for the graph to be displayed.

See [“Selecting storage objects and starting the data collection for historical statistics”](#) on page 239.

When you create a graph, you select which statistical parameters to use.

See [“Statistics overview”](#) on page 229.

To create a graph

- 1 Right-click the storage object you want to graph, and select **Graph Historical Statistics Information** from the context menu that appears.

A graphing window appears, but no graph is generated until you select a statistical parameter.

- 2 Choose **View** from the menu bar at the top of the screen, and then click one of the statistical parameters from the **View** menu.

Once you click a parameter, the **View** menu closes, and the graph is generated.

See [“Features of the graphing window”](#) on page 237.

Stopping historical data collection

You can stop the collection of historical data by using the Stop Historical Data Collection command. The historical data collection is also stopped when Storage Foundation for Windows is stopped or when the computer is restarted.

In those cases, you need to restart the historical data collection.

See [“Selecting storage objects and starting the data collection for historical statistics”](#) on page 239.

To stop the historical data collection

- 1 Select **Tools > Statistics > Stop Historical Data Collection** to bring up the Stop Historical Data Collection window, as shown in the sample screen that follows.
- 2 In the left pane of the window, highlight the storage objects for which you no longer want to collect historical data, and use the **Add** or **Add All** button to move the objects to the right pane. You can also move the objects to the right pane by double-clicking them or by dragging them to the right pane.
- 3 Click **OK** when you are satisfied with your selections.

Subdisk move, split, and join

This section describes the Subdisk Move, Split, and Join commands.

See [“Purpose of the subdisk commands”](#) on page 242.

See [“Subdisk context menu”](#) on page 243.

See [“Subdisk properties”](#) on page 243.

See [“Moving subdisks”](#) on page 244.

See [“Splitting subdisks”](#) on page 245.

See [“Joining subdisks”](#) on page 246.

Purpose of the subdisk commands

With the commands Move Subdisk, Split Subdisk, and Join Subdisk, you can move subdisks to other locations within the dynamic disk group, split subdisks, and join them back together. The flexibility of moving subdisks, splitting them, and joining them lets you make best use of your disk space.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Move Subdisk, Split Subdisk, and Join Subdisk commands.

Refer to the following topics to understand why you need to move or split your subdisks:

- See [“Load balancing”](#) on page 242.
- See [“Disk failure or removal”](#) on page 242.
- See [“Disk phaseout”](#) on page 243.

Load balancing

If disk activities are heavily concentrated on one or a small number of disks in the storage subsystem, it may create bottlenecks. You can use the Move Subdisk and possibly the Split Subdisk command to spread out disk accesses more evenly across all the disks.

If a disk has High or Critical I/O activity (shown by a yellow or red pie symbol), you may consider moving one or more of its subdisks to another disk that shows below-average I/O activity (shown by a blue pie symbol). The idea is to move enough activity to achieve balance. A careful study of the statistics for the disk with Critical activity may identify the best subdisks to move. You should move subdisks only when a disk has High or Critical I/O activity over a prolonged time and performance is affected. Moving a subdisk to another disk has an effect on I/O as well, but it should be compensated for by the other disk's having much lower I/O activity. You would need to look at the statistics after the subdisk move to see whether the move was effective in balancing the load.

Disk failure or removal

The Move Subdisk command can also be used to move the part of a volume that is on a failed or missing disk to a healthy one. Moving a subdisk for redundant volumes (mirrored or RAID-5) uses the redundant data to recreate the subdisk on the healthy disk. However, for nonredundant volumes (concatenated or striped),

the data cannot be recreated. In this scenario, if you do a subdisk move, you lose the data, which could be recovered if the disk is repaired. Thus, when you attempt to move a subdisk from a failed or missing disk that has nonredundant data, a dialog box appears that asks if you want to force the move. You may want to force the move if you don't need the data any more or you can no longer recover the data. By doing so, you retain the volume structure, but there is no guarantee that the data is recoverable.

Disk phaseout

When a disk starts getting intermittent I/O errors and shows signs of hardware fatigue, you can use the Move Subdisk command to move all its subdisks to healthier disks. The benefit of moving subdisks instead of copying the volumes is that you need only enough space on the receiving disks for the subdisks on the one failed disk, not for entire volumes that may span multiple disks. Another advantage is that there is no interruption in I/O.

Moving subdisks provides you a flexible means of making adjustments in your storage system while it is up and running.

Subdisk context menu

To access this menu, display the subdisk you want to work with by first clicking on its volume in the tree view and then clicking the Subdisks tab above the right pane. Once the subdisk is displayed, right-click the subdisk to bring up this menu.

Subdisk properties

Note: The Size, Disk offset, and Plex offset fields have pull-down lists that let you select the units you want the number to be displayed in (Sectors, KB, MB, GB, or TB).

The properties are as follows:

Name	In the screen captures showing subdisks in this chapter, you may have noticed that the subdisks on a disk are assigned numbers in a sequence, such as 1-01, 1-02, 1-03 or 12-01, 12-02, 12-03. The subdisk number is an internal number generated by the program. There is no particular significance to the one or two digits before the dash except that for any one disk, the first one or two numbers identifying its subdisks must be the same. The numbers after the dash represent a number for the subdisk within the disk. Thus, if the subdisk was the first one created for a particular disk, it would have a 01 after the dash, the second disk created would have a 02 after the dash, and so on.
Column	Column displays the column that the subdisk occupies in the volume. Striped and RAID-5 volumes have multiple columns.
Size	This field displays the size of the subdisk.
Disk offset	This is the offset where the subdisk starts on the disk. This is included in the Properties window because for a subdisk to be split, it must be contiguous to the subdisk next to it. By looking at the offsets, you can determine which subdisks are contiguous.
Plex offset	This is the offset of a subdisk within a plex. It shows the distance from the start of a plex to the start of a specific subdisk. The first subdisk within a plex has an offset of 0. The second subdisk is offset by the size of the first subdisk.

Moving subdisks

The purpose of moving a subdisk is discussed in another section.

See [“Purpose of the subdisk commands”](#) on page 242.

Note: You are not allowed to move a subdisk that is part of a boot or system volume.

Storage Foundation for Windows automatically checks the state of the data before beginning the subdisk move. If the data is degraded, the subdisk can still be moved. After moving, you need to reactivate the disk to clear the Degraded state.

Moving a subdisk from one disk to another involves a certain amount of time. The SmartMove feature can help reduce the time for this operation.

See [“SmartMove”](#) on page 93.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Move Subdisk command.

To move a subdisk

- 1 In the tree view, select the volume that contains the subdisk you want to move.
- 2 Identify the subdisk that you want to move, and display it in the right pane of the VEA console.

If you are concerned about load balancing, look for a disk with High or Critical I/O activity. It has a yellow or red pie symbol in the cell to the right of its name in the Statistics tabular view. You may also be aware that a disk has had difficulties because of minor disk errors that have occurred. Such a disk may display an Online (Errors) state from time to time.

To display the subdisk in the right pane, click on its volume in the tree view, and then click on the **Subdisks** tab in the right pane. You can use the subdisks' **Properties** command to further identify the characteristics of a volume's subdisks.

- 3 Right-click the subdisk you want to move to bring up its context menu.
- 4 Select **Move Subdisk** from the context menu.

The Move Subdisk window appears.

Choose either automatic or manual selection of the destination disk or disks for the subdisk to be moved to.

If you want to select the destination disk(s), click **Manually assign destination disks**. In the left panel of the window, highlight the disk(s) you want to assign as the destination disk or disks and use the **Add** or **Add All** button to move the disk(s) to the right pane of the window.

You may also check **Disable Track Alignment** to disable track alignment on the destination disks.

- 5 Click **OK** when you are satisfied with the changes.

Splitting subdisks

The Split Subdisk command lets you split a subdisk. After the subdisk is split, the resulting two subdisks reside on the same sections of the same disk. You can then use the Move Subdisk command to move one or both subdisks to other disks. You may want to split and move a subdisk to make the best use of your disk space. A split subdisk can be joined back together with the Join Subdisk command.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Split Subdisk command.

To split a subdisk

- 1 In the tree view, select the volume that contains the subdisks you want to split.
- 2 In the right pane, click the **Subdisks** tab to display the volume's subdisks.
- 3 Right-click on the subdisk you want to split to bring up its context menu.
- 4 Select **Split Subdisk** from the context menu.
- 5 In the window that comes up, specify the subdisk size for the first of the two subdisks.

The remaining space is assigned to the second subdisk. For example, if the original subdisk was 100 MB and you assigned 40 MB to the first subdisk, the other subdisk is 60 MB.

The drop-down list to the right of the **Size** and **First subdisk size** text entry boxes let you select sectors, KB, MB, GB, or TB.

- 6 Click **OK** to complete the command.

After a few moments, two subdisks appear in the right pane of the VEA console.

Joining subdisks

The Join Subdisk command lets you join subdisks. You can join two subdisks that had been one subdisk but that were split by the Split Subdisk command. You cannot split a subdisk and join it back to another subdisk that it was not split from. There also has to be room on the disk for the two subdisks.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Join Subdisk command.

To join subdisks

- 1 In the tree view, select the volume that contains the subdisks you want to join.
- 2 In the right pane, click the **Subdisks** tab to display the volume's subdisks.
- 3 Right-click on a subdisk you want to join to bring up its context menu.
- 4 Select **Join Subdisk** from the context menu.
- 5 In the window that appears, select the subdisks to be joined.
- 6 Click **OK** to complete the command.

Within a short time, the two subdisks in the right pane of the VEA GUI are replaced with one subdisk.

Command line commands for statistics

Statistics commands can be done through the command line interface.

See [“vxstat”](#) on page 681.

SmartMove

The performance of mirror operations and subdisk moves can be enhanced with the SmartMove feature. SmartMove helps reduce the resynchronization time required by mirror operations and subdisk moves. The resynchronization time is reduced by using the NTFS or ReFS file system metadata to resynchronize only selected regions. Operations that involve mirrors, like adding a mirror to a volume, off-host backup, and array migration, may have improved performance by using the SmartMove feature.

More information about SmartMove is available.

See [“SmartMove”](#) on page 93.

Note: SmartMove supports only NTFS and ReFS file systems.

FlashSnap

This chapter includes the following topics:

- [FlashSnap overview](#)
- [FlashSnap components](#)
- [FastResync](#)
- [Snapshot commands](#)
- [Dynamic Disk Group Split and Join](#)
- [CLI FlashSnap commands](#)
- [Fast File Resync](#)
- [Volume Shadow Copy Service \(VSS\)](#)
- [Using the VSS snapshot wizards with Microsoft Exchange](#)
- [Using the VSS snapshot wizards with Enterprise Vault](#)
- [Using the VSS snapshot wizards with Microsoft SQL](#)
- [Copy on Write \(COW\)](#)
- [Using the VSS COW snapshot wizards with Microsoft Exchange](#)
- [Using the VSS COW snapshot wizards with Microsoft SQL](#)

FlashSnap overview

Veritas FlashSnap is a multi-step process that lets you create independently addressable snapshot volumes that are copies or mirrors of the volumes on your server. These snapshots can be used for backup, application testing, or reporting and analysis. One of the main advantages of using FlashSnap is that the snapshot

volumes can be easily moved to another server. Thus, backup or other processing can be performed on another server without affecting the performance of applications on the applications server.

With Storage Foundation (SFW), you can also create simultaneous, multiple split-mirror snapshots. These snapshots can be done either through the VEA GUI or through the `vxsnap` CLI command.

Note: FlashSnap is available only with the Enterprise edition. See the "Licensing" section in the *Veritas InfoScale Installation and Upgrade Guide* for details.

SFW supports both traditional full snapshots or Microsoft Copy on Write (COW) snapshots. The VEA GUI or the `vxsnap` CLI utility can be used to take COW snapshots. Microsoft COW snapshots are faster than full snapshots and take up less storage space because COW snapshots use a shadow copy that is differential. The initial state of the original volume can be reconstructed using the blocks in the differences area and unchanged blocks in the original volume.

SFW provides support for taking snapshots of Microsoft SQL databases and Microsoft Exchange mailbox databases. FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all the volumes that are associated with an SQL or Exchange database without taking the database offline.

The VSS Snapshot wizard and the `vxsnap` CLI command both use VSS and either can be used to take snapshots of storage groups. The VSS process allows the storage group's databases to be quiesced before the snapshot operation and reactivates them immediately after. The quiescing at the storage group level allows for Microsoft supported and guaranteed persistent snapshots of the data. Snapshots of the storage groups can be reattached and resynchronized to match the current state of the storage group with the VSS Snapback wizard or the `vxsnap reattach` CLI command.

Refer to the *Storage Foundation and High Availability Solutions High Availability and Disaster Recovery Solutions Guide for Microsoft Exchange* for detailed procedures on how to use the VSS Snapshot wizard or the `vxsnap` utility with Microsoft Exchange to perform and to implement recovery procedures. This guide is included in PDF format with the product.

See "[vxsnap](#)" on page 710.

See "[Using the VSS snapshot wizards with Microsoft Exchange](#)" on page 295.

SFW also provides a VSS Snapshot Scheduler Wizard that enables you to set up a schedule for automating the snapback process.

At the scheduled time for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then snapshotted again. A scheduler service, VxSchedService.exe, that runs in the background, maintains the schedule.

See [“Using the VSS Snapshot Scheduler wizard”](#) on page 301.

Fast File Resync is a feature that uses Veritas FastResync technology to resynchronize selected files in a snapshotted volume to the original volume. It is designed for database applications using volumes with multiple database files and can be used to restore one or more corrupt files.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support FlashSnap or SFW VSS-related commands.

FlashSnap components

FlashSnap is a multi-step process.

It includes the following components:

- **FastResync**
FastResync (FR) supports resynchronizing of mirrors by copying only changes for the temporarily split mirror by using FR logging. This reduces the time it takes to rejoin a split mirror to the mirror set and also reduces the server CPU cycles that are needed to complete the resynchronization.
The Fast File Resync command provides the capability to restore a single file from a snapshot back to the original volume.
See [“Fast File Resync”](#) on page 287.
- **Snapshot commands**
The snapshot commands are used to create the mirrored volumes that are useful for backup or other resource-intensive processing purposes.
These features are necessary for the FlashSnap procedure, but they can also be used for other, more general purposes.

Note: Some snapshot commands, such as Prepare, may take a significant amount of time. These mirror-related operations can be tuned for better performance with the Configuration Task Performance Tuning settings.

See [“Configuration task performance tuning”](#) on page 95.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot commands.

- Dynamic Disk Group Split and Join

Dynamic Disk Group Split and Join supports the ability to split a dynamic disk group into two disk groups so that the newly formed disk group can be moved to another server. This allows splitting a mirror for backup and having a separate server handle the backup. After the backup is completed, the split-off disk group is moved back to the original server and joined to its former disk group, and the mirror is reassociated with its mirror set and resynchronized. Dynamic Disk Group Split and Join also can be done on the same server for same-host backup or for reorganizing the disk groups on the server.

The Dynamic Disk Group Split and Join commands are necessary for FlashSnap, but they can also be used for other, more general purposes.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Dynamic Disk Group Split and Join commands.

Summary of the FlashSnap procedure

You can use FlashSnap with a single server or with two servers (using a remote server for processing the backup or other resource-intensive activities).

This section includes a summary for both situations. These steps can be implemented through the VEA GUI or through the command line.

Note: The same host procedure can be used for quick recovery snapshots to recover a data.

Off-host FlashSnap procedure (Two servers)

The following is the off-host FlashSnap procedure.

To use the off-host FlashSnap procedures with two servers

- 1 Use the **Prepare** command to attach a snapshot mirror to a volume on the working server.

You need to perform this step only once for the initial setup.

The **Prepare** command replaces the **Snap Start** command in the GUI. Both `prepare` and `start` keywords are available in the CLI, however `prepare` is the recommended keyword.

- 2 Use the **Snap Shot** command to automatically detach the snapshot mirror and create a new volume that is associated with the snapshot mirror.

This new volume can be used for backup or other purposes. The snapshot process typically takes less than 20 seconds.

- 3 Split the new volume from its dynamic disk group with the **Split Dynamic Disk Group** command, creating a new dynamic disk group on the server.

- 4 Deport the new dynamic disk group from the original server.

When you use the command line, a separate deport command is not necessary. It is needed only for the VEA GUI command.

- 5 Rescan on the second server so that the computer recognizes the split-off dynamic disk group.

- 6 Import the dynamic disk group on the second server where the backup or other off-host processing occurs.

- 7 Perform the backup or other off-host processing activity, using the snapshot volume.

To backup the server, NetBackup or Backup Exec is recommended for the backup software.

- 8 Deport the dynamic disk group from the second server.

If you do not have a SAN where you can access the deported dynamic disk group from the original server, physically move the split-off disk group's disks back to the original server.

- 9 Rescan on the original server, so that the server recognizes the split-off dynamic disk group.

You do not need to import the dynamic disk group back to the original server.

- 10 Use the **Join Dynamic Disk Group** command to join the split-off dynamic disk group back to its original dynamic disk group.
- 11 Join the mirrored volume back to its original volume, using the **Snap Back** command.

The mirrored volume is now resynched to the data source and is continually updated. Because FastResync (FR) is used in resynchronizing the mirrors, the time for the resynchronization is greatly reduced.

The next time the sequence of commands is run again, the snapshot mirror is ready for the snapshot step. The **Prepare** step, step 1, does not have to be done again.

Same-host FlashSnap procedure (One server)

Steps 1 and 2 are the same as in the [Same-host FlashSnap procedure \(One server\)](#) section:

- 1) Using Snap Prepare to attach a snapshot mirror to the original volume
- 2) Detaching the snapshot mirror and creating a new volume that is associated with it.

Because no remote host is involved, you do not need to use the Disk Group Split and Join commands. The backup or other resource-intensive processing is done on the same server.

In both the off-host and same-host procedures, the final step is to join the mirrored volume back to its original volume, using the Snap Back command. In this step, the time for synchronizing the mirrors is greatly reduced.

FlashSnap on the same host can also be used for quick recovery snapshots to recover data. Quick Recovery is the process of using on-host point-in-time copies of production data and a transaction log to recover a database that has been corrupted or that has missing data. If a database becomes corrupted, for example, you could reload the original data from the most current snapshot, and then use the transaction log to bring the database current to the point before the corruption.

For a general discussion of Quick Recovery, see the Quick Recovery section of the *Storage Foundation and High Availability Solutions Solution Guide*. For detailed information, refer to the Quick Recovery Solutions guide for the respective application.

FastResync

FastResync (FR) provides enhancements to Storage Foundation.

These include the following:

- Resynchronization
Optimizes the resynchronization of a mirror that is detached from and then returned to a volume.
- Reconnection
Extends the snapshot model by providing the ability to refresh and reuse a snapshot, rather than discarding it.

FR increases the efficiency of the Storage Foundation snapshot mechanism and improves the performance of operations, such as backup and decision support. These operations are organized methods of collecting data and are designed to facilitate decision making). Typically, these operations require that a volume is quiescent and that they are not impeded by updates to the volume by other activities in the system. FR reduces the amount of time that a volume must be out of service during such operations.

You can enable or disable FR on a per-volume basis. When you enable FR for a volume, a Disk Change Object (DCO) volume is created to track the regions on a volume that are changed while a mirror is detached.

Note: Only a limited subset of operations is permitted on a DCO volume. You can perform a subdisk move, or add and remove a mirror.

FR is automatically enabled for a volume when **prepare the volume for snapshot operation** is performed on the volume.

The volume can be prepared through the command line interface `vxassist prepare` command or through the VEA GUI Prepare command.

Note: The Prepare command replaces the Snap Start command in the VEA GUI.

Both `prepare` and `start` keywords are available in the CLI, however `prepare` is the recommended keyword.

FastResync limitations

FastResync cannot be used with software RAID-5 volumes.

Resynchronization

FastResync (FR) keeps a track of the datastore updates that a detached mirror missed. When a mirror returns to service, FR resynchronizes the mirror with the volume by applying only those updates that the mirror missed while it was detached.

This generally takes less time than the traditional method of copying the entire datastore to the returning mirror.

If you want FR to resynchronize a mirror, FR must be enabled for the volume before the mirror is detached and continue to be enabled until after the mirror is reattached and resynchronized. However, if you use any of the snapshot commands, FR is automatically turned on.

FR does not alter the traditional mirror failure and repair administrative model. The only visible effect is that typical mirror repair operations conclude more quickly.

Although both FR and dirty region logging (DRL) keep track of regions on a volume where the mirrors are not synchronized, they perform different functions. FR keeps track of the datastore updates that a detached mirror missed, while DRL keeps track of whether a write to a mirrored volume has been completed on all mirrors. The write region on the volume is considered "dirty" because the mirrors are out of sync until the write to all mirrors is completed. DRL is used to resynchronize mirrors following a system crash.

See ["Dirty region logging \(DRL\) and RAID-5 logging"](#) on page 222.

Reconnection

Reconnection extends the snapshot model by providing the ability to refresh and reuse a snapshot mirror, rather than discarding it. FR is automatically enabled when you take a snapshot of a volume. FR retains the association between the snapshot mirror and the volume when the snapshot mirror is detached from the volume.

See ["Resynchronization"](#) on page 254.

FR reattaches and resynchronizes the snapshot mirror with the volume when the Snap Back command is issued, either through the `vxassist` command line command or the GUI.

See ["vxassist"](#) on page 638.

See ["Snapshot commands"](#) on page 259.

Enabling and disabling FastResync

You can enable or disable FastResync (FR) on a per-volume basis. For FR to be active when a mirror is reattached and resynchronized with its associated volume, FR must be enabled before the mirror is detached and remain enabled until the mirror has been reattached and resynchronized with the volume.

The procedure for enabling FR from the GUI is described below.

You can also enable and disable FR from the command line.

See “vxvol” on page 581.

Note: You cannot enable FR on RAID-5 volumes.

Note: If you use snapshot commands, you do not need to use the following steps, because FR is automatically enabled for snapshot commands. These steps are needed only when you want to enable FR on a volume that is not used with any snapshot commands.

Enabling FastResync

The following describes how to enable FR.

To enable FastResync for a volume through the VEA GUI

- 1 Select the mirrored volume that you want to enable FR for.
The volume context menu is displayed.
- 2 Select **FastResync > Add**.
A dialog box is displayed, asking you to confirm adding FR for the volume.
- 3 Click **Yes**.
FR is now enabled for the volume.

Disabling FastResync

The following describes how to disable FR.

To disable FastResync for a volume through the VEA GUI

- 1 Select the mirrored volume that you want to disable FR for.
The volume context menu is displayed.
- 2 Select **FastResync > Remove**.
A dialog box is displayed asking you to confirm removal of FR.
- 3 Click **OK**.
FR is now disabled for the volume.

Disk Change Object (DCO) volume

The DCO volume is created when you enable FR or when a snapshot operation is started. The DCO volume keeps track of the changes that are made to a volume while a mirror is detached.

The DCO volume is not visible in the tree view in the left pane of the VEA console. It is visible in the Disk View (when volume details is not selected).

You can perform the following on a DCO volume:

- Add or remove a mirror — you can add a mirror to the DCO volume and remove a mirror from it.
- Move a subdisk — you can move the DCO volume and the mirrored subdisk that is associated with it to another disk with the Subdisk Move command.

Add a mirror to a DCO volume

You might want to add a mirror to a DCO volume to have a backup copy of the DCO volume.

To add a mirror to a DCO volume

- 1 Select the volume whose DCO volume you want to add a mirror to.
- 2 From the volume's context menu, select **Disk View** to get a separate Disk View window. You can also select the **Disk View** tab to have the Disk View display in the right pane.
- 3 Right-click the DCO volume to bring up a context menu, as shown in the sample screen.
- 4 Select **Mirror > Add**.
The Add Mirror window is displayed.
- 5 If you want to add the mirror to a specific disk, click the radio button **Manually select disks**, select the desired disk, and then click **Add** to move the disk to the right pane. If you want the disk to be selected automatically, leave the default setting of **Auto assign destination disks** selected.
- 6 Click **OK** to add the mirror.

If you add mirrors to the DCO volume of a snapshot volume, these mirrors are deleted when you snap back the snapshot volume.

The process for adding a mirror to a volume that is not a DCO volume is different.

See [“Add a mirror to a volume”](#) on page 208.

Remove a mirror from a DCO volume

The following is the procedure to remove a mirror from a DCO volume.

To remove a mirror from a DCO volume

- 1 Select the volume whose DCO volume you want to remove a mirror from.
- 2 Select the **Disk View** tab in the right-hand pane, and right-click on the DCO volume.
A context menu is displayed.
- 3 Select **Mirror > Remove**.
A dialog box is displayed, asking you to select the mirror to remove.
- 4 Select the mirror.
Click **Add** to move it to the right pane.
- 5 Click **OK**.
The mirror is removed from the DCO volume.

Move the DCO volume subdisk

You might want to move a DCO volume and the associated mirrored subdisk to another disk if something happens to the original disk, or you may prefer to have this subdisk on another disk to split the disk group differently. The DCO volume and the associated mirrored subdisk have to be moved to a disk in the same disk group.

To move a DCO volume subdisk

- 1 Under the Volumes folder, select the volume whose DCO volume subdisk you want to move.
- 2 Right-click in the empty space on the **Disk View** tab in the right-hand pane.
A context menu is displayed.
- 3 Select **Vol Details**.
The display changes. The **Disk View** tab now displays the subdisk names for the volumes.
- 4 Right-click on the subdisk containing the DCO volume.
A context menu is displayed.
- 5 Select **Move Subdisk**.
The Subdisk Move window appears.
- 6 Click **Manually assign destination disks**.

- 7 Highlight the disk you want to move the mirrored volume subdisk with the DCO volume to, and click **Add**. The selected disk is moved to the right pane of the window.
- 8 Click **OK** to complete the command.
The mirrored volume subdisk with the DCO volume is moved to the other disk.
More information about subdisk move is available.
See [“Subdisk move, split, and join”](#) on page 241.

Snapshot commands

You can use the Veritas Enterprise Administrator (VEA) GUI to take a snapshot of a volume.

The snapshot feature can also be accessed through the command line interface.

See [“CLI FlashSnap commands”](#) on page 286.

You can take a snapshot of mirrored or unmirrored volumes, but not of RAID-5 volumes.

The Snapshot feature enables quick and efficient online backup of volumes, with minimum disruption to the user. Snapshot automatically enables FastResync (FR) and creates a snapshot image of a volume, which can be detached from the volume. The detached snapshot image can be used for backup or other purposes, while the original volume remains in service.

The snapshot can then be quickly reattached and resynchronized with the volume through the FR reconnection capability.

See [“Reconnection”](#) on page 255.

The snapshot volume can also be permanently disconnected from the original volume.

You can attach more than one snapshot plex to a volume by issuing successive Prepare commands on the volume. This enables you to take multiple snapshots of a single volume.

Note: The Prepare command replaces the Snap Start command in the VEA GUI.

You can attach a snapshot to either the original volume or the snapshot volume.

The GUI supports the following snapshot commands. You can access these commands through the volume context menu:

- Prepare

Creates the snapshot mirror and attaches it to the original volume.

- Snap Shot
Detaches the snapshot mirror from the original volume and creates a snapshot volume that is associated with the snapshot mirror.
- Snap Back
Reattaches the snapshot mirror to the original volume.
- Snap Clear
Permanently removes the snapshot mirror from the original volume.
- Snap Abort
Aborts the snapshot operation after a Prepare or Snap Back command is issued. Snap Abort permanently removes the snapshot mirror from the volume and releases its space.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot commands (Prepare, Snap Shot, Snap Back, Snap Clear, or Snap Abort).

Using the snapshot procedure

In summary, the snapshot procedure consists of a few steps.

These include the following:

- Attaching a snapshot mirror to a volume, using Prepare.
- Detaching the snapshot mirror and creating the snapshot volume that is associated with the snapshot mirror.
- Either reattaching the snapshot to the volume after you are done using it, or permanently disconnecting it from the volume. Snap Back reattaches the snapshot mirror; Snap Clear permanently disconnects the snapshot mirror from the volume, removing the association between the snapshot volume and the original volume.

The Prepare procedure takes considerable time because it involves creating a mirror, but it has to be done only the first time you perform the snap commands sequence. After that, you use the commands Snap Shot to break off the mirror for backup (or other purposes) and then Snap Back to reattach the snapshot mirror to the original volume. Both of those are very quick procedures (typically less than a minute for each one).

Snap Abort aborts the snapshot operation after the Prepare or Snap Back command is issued.

Prepare

Prepare creates a snapshot mirror or plex, which is attached to and synchronized with a volume. Alternatively, if you apply the command to a volume that already has one or more normal mirrors, you can designate an existing mirror to be used for the snapshot mirror. The advantage of selecting an existing mirror is that it saves time, since it is not necessary to resynchronize the mirror to the volume.

Note: The Prepare command replaces the Snap Start command in the VEA GUI.

The mirror synchronization process can take a while, but it does not interfere with use of the volume. If the prepare the volume for snapshot process fails, the snapshot mirror is deleted if it was created from scratch, and its space is released. If you selected a normal mirror to be used for the snapshot mirror, that mirror reverts to its normal state if the prepare the volume for snapshot process fails.

When the prepare the volume for snapshot process is complete, the status of the snapshot mirror displays as Snap Ready on the Mirrors tab in the right pane of the VEA GUI. The snapshot mirror can be associated with a snapshot volume by using the Snap Shot command. Once the snapshot mirror is created, it continues to be updated until it is detached.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Prepare or Snap Start commands.

To create a snapshot mirror

- 1 Right-click on the volume that you want to take a snapshot of.
A context menu is displayed.
- 2 Select **Snap > Prepare**.
The Prepare volume for FlashSnap wizard welcome screen appears.
Click **Next** to continue.
- 3 The screen that appears depends on whether you already have a mirrored volume or not. If you already have a mirrored volume, the snapshot mirror continues to be updated until it is detached by using the **Snap Shot** command.
The various screens are as follows:
 - Mirrored volume: If you have a mirrored volume, a screen appears to let you select an existing mirror to be used for the snapshot mirror.

- If you have a mirrored volume and there is also a disk available on your system to create an additional mirror, the screen lets you choose either to use an existing mirror for the snapshot or to have a new mirror created.
- If you have a mirrored volume and there is no disk available for creating a new snapshot mirror, the screen lets you select from existing mirrors in the volume.

If you select an existing mirror, click **Next** to continue to the summary screen and click **Finish** to complete the **Prepare** command.

If you do not select an existing mirror, click **Next** to continue and follow the instructions for an unmirrored volume.

- **Unmirrored volume:** If you have an unmirrored volume or you have not selected an existing mirror to use for the snapshot mirror, select the disk to be used for the snapshot mirror from the window for disk selection. The default setting is to have the program automatically select the disks where the mirror is created. Alternatively, you can specify the disks that can be used to create the snapshot mirror by clicking the **Manually select disks** radio button. If you select the manual setting, use the **Add** or **Add All** option to move the selected disks to the right pane of the window. The **Remove** and **Remove All** options let you move selected disks back to the left pane. You may also check **Disable Track Alignment** to disable track alignment on the snapshot mirror volume. Click **Next** to continue to specify attributes.

- **Specify attributes**

On this screen select one of the following volume layout types:

- **Concatenated**
- **Striped**

If you create a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.

For a concatenated or striped volume, you may also specify to mirror across disks by the following:

- **Port**
- **Target**
- **Enclosure**
- **Channel**

The operation to prepare a volume for a snapshot fails if the appropriate resources are not available to support the selected attributes to mirror across disks.

After the Prepare command completes, a new snapshot mirror is attached to the volume. See the sample screen below. In that screen, the volume Flash has a snapshot mirror attached to it.

The new mirror is added to the **Mirrors** tab for the volume. In the sample screen, the mirror is identified as a snapshot mirror and has the Snapshot icon. After the snapshot mirror is synchronized with the volume, its status becomes Snap Ready.

It is important to make sure that the snapshot mirror (or snap plex) has completed its resynchronization and displays the status of Snap Ready before continuing with the Snap Shot command or doing any other operations on the snapshot mirror. Also, if you shut down the server or deport the disk group containing the volume being prepared for a snapshot before resynchronization completes, the snapshot mirror is deleted when the disk group with the original volume comes online again.

The DCO (Disk Change Object) volume is created to track the regions on a volume that are changed while a mirror is detached.

See [“Disk Change Object \(DCO\) volume”](#) on page 256.

The DCO volume is not included in the tree view of the VEA GUI. To view the DCO volume, you must use the Disk View. To access the Disk View, click the Disk View tab in the right pane or select Disk View from a disk's or volume's context menu.

The sample Disk View screen that follows shows the DCO log that the Prepare command creates.

Note: The Break Mirror and Remove Mirror commands do not work with the snapshot mirror.

Snap shot

Snap Shot detaches the snapshot mirror and creates a new volume that is associated with that mirror. This process takes a relatively short time, typically less than a minute, during which users can be requested not to access the original volume. After the snapshot volume is created, it can be used for backup or other purposes. Snap Shot can also make snapshots of more than one volume at the same time.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Snap Shot command.

To take a snapshot

- 1 Right-click on one of the volumes with a snapshot mirror. A context menu is displayed.
- 2 Select **Snap > Snap Shot**.
- 3 The **Snap Shot Volume** Wizard Welcome screen appears. Click **Next** to continue.
- 4 In the **Select Volumes** window that appears, select the volumes that you want to take a snapshot of.

Highlight the volumes in the **Available volumes** pane and click **Add**. Your selected volumes should appear in the **Selected volumes** pane.

You have the option of taking the snapshot with Volume Shadow Copy Service (VSS). The VSS snapshot method lets you take snapshots of VSS-aware applications, such as Microsoft SQL Server, while the application files are open. When VSS-aware applications do not exist, the snapshot is taken with the SFW FlashSnap method (VM method).

See [“Volume Shadow Copy Service \(VSS\)”](#) on page 289.

When the VSS snapshot method is used, snapshots can only be taken of read or write volumes and the resulting VSS snapshot is read only. The V snapshot method may take snapshots of read or write or read-only volumes and the resulting snapshots match the access mode of the original volume. However, when the VM snapshot method is used the **Snap Shot Volume** Wizard presents a **Force Read Only** check box to let you specify the resulting snapshot to have read only access.

You can set the snapshot process to use the VSS or the VM snapshot method by using Systems Settings.

See [“System settings”](#) on page 89.

Click **Next** after selecting your volumes.

- 5 In the **Change Attributes** window that appears, you may edit the following attributes of the snapshot:
 - Snap Volume Label

Note: The default Snap Volume Label begins with “SnapV” followed by a 12 character alphanumeric string. The 12 character alphanumeric string is a coded hexadecimal timestamp based upon the time that the snapshot was taken. For example, a default Snap Volume Label would be “SnapV4A7337AE038A”.

Note: The Snap Volume Label cannot be changed for read-only snapshots.

- Drive Letter
- Plex

After editing the attributes, click **Next** to continue.

6 Click **Finish** in the final window to begin the Snap Shot process.

The snapshot mirror is detached from each of the original volumes, and new volumes are created that are associated with the snapshot mirrors. This process usually takes less than a minute.

The snapshot mirrors are no longer displayed on the **Mirrors** tab of the original volumes. The new snapshot volumes are displayed under the Volumes folder in the tree view. The program assigns them the next available drive letters (drive letters F and L in the tree view of the following sample screen).

You can now use the snapshot volume for backup or other purposes.

Snap back

Snap Back reattaches a detached snapshot mirror to its original volume.

After you are done using the snapshot, you can reattach it to the volume, where it is resynchronized to the datastore and continually updated. You can then get a current copy of the original volume, with minimal disruption to users, by issuing Snap Shot again.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Snap Back command.

To reattach a snapshot mirror to the original volume

- 1 Right-click the snapshot volume whose snapshot mirror you want to reattach to its original volume.

A context menu is displayed.

- 2 Select **Snap > Snap Back**.

A dialog box appears with two options for resynching the snapshot.

- 3 Specify whether the snapshot volume is to be resynchronized to match the original volume or the original volume is to be resynchronized to the snapshot volume, and click **OK**.

One situation where you might want to resynchronize using the snapshot volume rather than the original volume is when something has happened to the original volume so that its data integrity is no longer sound.

The middle of the screen displays a further explanation of the selected option.

The snapshot mirror is reattached to the original volume, and the snapshot volume is deleted if it is the last plex of the volume. (If the snapshot volume has other associated mirrors, it is not deleted.)

As shown in the following figure, the snapshot volume, identified as L in the previous figure, is removed from the Volumes folder in the tree view, and the snapshot mirror is now shown on the **Mirrors** tab for the original volume. After the resynchronization, the original volume is now in a similar situation as it was after the first **Prepare** command, but any changes to it have been updated on the snapshot mirror.

Now it is possible to do another **Snap Shot** command on the volume and repeat the process. If you resynchronize to the original volume, the resynchronization should take very little time because that volume is constantly updated.

Only the original synchronizing of the mirror in the original **Prepare** command takes the normal amount of time to create the snapshot mirror. Once you have completed the original **Prepare** command, the repeat procedure for **Snap Shot** and **Snap Back** are very quick. You have to do the **Prepare** command only once.

Note that any mirrors that are added to the DCO volume of a snapshot volume are deleted when you snap back the snapshot volume.

Snap clear

Snap Clear permanently removes the link between a snapshot volume and its original volume. Snap Clear can be specified after the Snap Shot command has been issued. After a Snap Clear, the snapshot volume becomes an ordinary volume.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Snap Clear command.

To change a snapshot volume to an ordinary volume

- 1 Right-click on the snapshot volume that you want to change to an ordinary volume.
- 2 Selecting the snapshot volume and after the **Snap Clear** command is completed, the volume is no longer linked to its original volume.
- 3 Select **Snap > Snap Clear**.
A message box is displayed asking you to confirm the **Snap Clear** operation for the specified volume.
- 4 Click **Yes**.
The snapshot mirror becomes an ordinary volume. The original volume goes back to the state that it was in before the **Prepare** command.

Snap abort

Snap Abort aborts the changes that the Prepare or Snap Back commands make. In both these commands, a snapshot mirror plex is attached to a volume. Snap Abort either deletes this snapshot mirror plex or converts it to an ordinary mirror. In cases where the deleted snap plex is the last snap plex and the resulting volume is simple or striped, the Snap Abort command deletes the DCO log volume also. The command cannot be done directly after a Snap Shot command.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Snap Abort command.

To delete a snapshot mirror or change it into an ordinary mirror

- 1 Right-click the volume on which you want to abort the **Prepare** operation.
- 2 Select **Snap > Snap Abort** from the context menu that comes up.

- 3 Select the options in the **Snap Abort Volume** window that appears.

If you have only one mirror, you can leave the default automatic selection setting. If there are multiple mirrors, you can manually select the mirrors to snap abort. If you click the check box **Convert into a normal mirror**, the selected mirror or mirrors become normal mirrors.

Click **OK** when you are satisfied with your selections.

- 4 Respond to the message box asking you to confirm the **Snap Abort** operation for the specified volume. The snapshot mirror is deleted or converted. The DCO volumes are removed.

Using the Snapshot Scheduler wizard

You can use the Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of taking snapshots. At the scheduled time, the specified volumes are automatically snapshotted. Once configured and applied, the schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

Note: You can find the information about the operations that the schedule performs by expanding the nodes under the Scheduled Tasks node in the tree-view of the GUI. Expand the appropriate application node to locate the schedules for that application. Right-click the node and select **Job History** to view information about the operations performed by the schedule.

Note: The Snapshot Scheduler only supports VSS snapshots (read-only snapshots).

Note: You must prepare the snapshot mirror with the `Prepare` command before running the Snapshot Scheduler wizard.

To create a schedule for snapshots

- 1 Right-click a volume in the tree-view.
- 2 Click **Snap > Schedule SnapShot** in the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.

- 4 In the **Select Volumes** window that appears, select the volumes that you want to take a snapshot of.

Highlight the volumes in the **Available volumes** pane and click **Add**. Your selected volumes should appear in the **Selected volumes** pane.

Specify the snapshot set parameters as appropriate and then click **Next**

Directory The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory that is shown on the screen.

You can retain this location for non-clustered environments. For clustered environments, enter the Universal Naming Convention (UNC) path to a location on shared storage. For example, `\\ServerName\SharedFolderName`.

Snapshot set Enter a name for the snapshot set, for example, `billing` or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet.vol.timestamp".

- 5 In the **Change Attributes** window that appears, you may edit the following attributes as appropriate. Click **Next** to continue.
 - Drive Letter
 - Plex

6 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

Name of this schedule	<p>Enter a unique name for the snapshot set schedule.</p> <p>This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of a prefix, "Generic", and a numbered suffix that increments with each schedule.</p>
Description of this schedule	<p>Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.</p>
Start Time	<p>The time of the day to begin taking snapshots.</p>
End Time	<p>The time of day to end taking snapshots.</p> <p>If a snapshot is in progress it is completed but a new one is not started after the end time.</p>
Schedule takes effect on	<p>The date on which the specified schedule takes effect. The default is the current date.</p>
Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

7 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

8 Click **Next**.

9 Review the specifications of the snapshot set and click **Finish**.

Displaying the status of the scheduled snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1** From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2** Expand the system node and the **Storage Agent** node.
- 3** Click **Scheduled Tasks**.

- 4 Expand the **Generic** node. The scheduled snapshots are listed on the right pane of the VEA.
All scheduled snapshots that are not associated with an Application are displayed under the Generic node.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Generic** and click **All Job History**
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the dialog box, view the schedule information.
You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a snapshot

The following is the procedure to delete a schedule.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the GUI console.

To delete a schedule with the GUI

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the **Storage Agent** node.
- 3 Click **Scheduled Tasks**.
- 4 Expand the Generic node for which you scheduled the snapshot. The scheduled snapshots are listed on the right pane of the VEA.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Dynamic Disk Group Split and Join

This section describes Dynamic Disk Group Split and Join.

See [“About Dynamic Disk Group Split and Join”](#) on page 273.

See [“Dynamic disk group split ”](#) on page 274.

See [“Recovery for the split command”](#) on page 277.

See [“Dynamic disk group join”](#) on page 279.

See [“Using Dynamic Disk Group Split and Join with a cluster on shared storage”](#) on page 282.

See [“Limitations when using dynamic disk group split and join with Volume Replicator”](#) on page 285.

See [“Dynamic Disk Group Split and Join troubleshooting tips”](#) on page 285.

About Dynamic Disk Group Split and Join

Dynamic Disk Group Split and Join refers to two related commands—Split Dynamic Disk Group and Join Dynamic Disk Group. Split Dynamic Disk Group splits a dynamic disk group into two dynamic disk groups. Join Dynamic Disk Group joins two dynamic disk groups into one merged disk group. You can join two disk groups that were originally split apart with the Split Dynamic Disk Group command, but you can also join two dynamic disk groups that started out as separate disk groups.

Dynamic Disk Group Split and Join can be used for the FlashSnap process and to reorganize dynamic disk groups.

It can be implemented through the VEA console or through the command line by using the `vxdg split` and `vxdg join` commands.

With the Split Dynamic Disk Group command, you can take some but not all disks from one dynamic disk group to another. The source dynamic disk group retains its identity as the original, while the other dynamic disk group, called the target disk group, becomes a new dynamic disk group. After the split operation, the target dynamic disk group is in the Imported state if you used the GUI to implement the command. If you used the command line to do the split, the target dynamic disk group is by default in the Deported state (though you can use the `-i` switch to have it remain in the Imported state). In both the GUI and the command line, the source dynamic disk group continues to remain online after the split operation.

Primary dynamic disk groups cannot be split because primary dynamic disk groups usually contain the computer’s boot and system disks. Also, only healthy dynamic volumes can be transferred in the split operation.

You can upgrade disk groups through the GUI.

See [“Upgrading the dynamic disk group version”](#) on page 150.

You can also upgrade disk groups through the command line.

See [“vxdg upgrade”](#) on page 610.

The Split Dynamic Disk Group command works on the premise that the split-off disk group should contain all disks that are needed to make the volumes in the new disk group complete. If the disks that you select to split the disk group result in

incomplete volumes, the logic that is built into the command adds the remaining disk or disks needed to split the disk group with complete volumes. A screen is presented to you that lists the disks that are needed for the split. You can decide at that time whether you want to go ahead with the dynamic disk group split.

Thus, you are not able to split a dynamic disk group into two disk groups if any of the volumes that are allocated to either split disk group are incomplete.

The Join Dynamic Disk Group command lets you combine dynamic disk groups. It does not require that the dynamic disk groups to be joined were previously split. When you join two dynamic disk groups, the dynamic disk group you designate as the source becomes the one that loses its identity in the merge. The target dynamic disk group is the one that remains after the join.

With the join command, all the disks of the source dynamic disk group are joined with all the disks of the target dynamic disk group. Storage Foundation allows a partial join—that is, volumes in the source disk group can have missing or nonfunctional disks and the join command can still succeed. In this situation, you must click the **Allow Partial Join** check box to indicate that you want a partial join to take place. When the missing disk or disks come back to an online state, then you need to do another Join Dynamic Disk Group command to add the missing disk or disks to the target disk group.

Warning: If you have a partial join in place, you must wait until the missing disks or disks are joined to the target disk group before making the following configuration changes to either the source or target disk group: replace or evacuate disk, split or move a subdisk, extend a volume, add a mirror, and clear hot relocation. Doing so before all disks are joined can result in data loss.

The procedures in this section focus on the CLI commands that are related to Dynamic Disk Group Split and Join.

You can also perform Dynamic Disk Group Split and Join through the command line.

See [“vxdg”](#) on page 593.

Note: Microsoft Disk Management Disk Groups do not support Dynamic Disk Group Split and Join.

Dynamic disk group split

After a dynamic disk group is split through the CLI command, the target disk group is in an Imported state. If you want the dynamic disk group to be in a Deported

state, use the CLI version of the command, `vxdg split`, in which the default state of the target disk group is `Deported`. However, you can also use the `-i` switch with `vxdg split` to have the target disk group be `Imported`.

With the dynamic disk group split command, if the source disk group is a cluster disk group or a disk group with private dynamic disk group protection, the resulting target disk group becomes the same disk group type.

With the CLI disk group split command, these disk group types need additional parameters to ensure that they retain their respective disk group type.

See “[vxdg split](#)” on page 612.

Note: Microsoft Disk Management Disk Groups do not support dynamic disk group split.

To split a dynamic disk group

- 1 Make sure that the volumes to be split are not in use and that they are healthy. Similarly, make sure that any disks that are split do not have a `Missing` status.

You can use the `Disk View` to look at the disk group’s volumes and subdisks, so you can see which disks are appropriate to split. The disks in the new split-off disk group have to contain complete volumes.

If a volume is in use, an error message comes up, asking if you want to force the split operation.

See “[Error message about forcing a split or join operation](#)” on page 286.

- 2 Right-click a dynamic disk group, disk, or volume in the tree view of the `VEA` console, and click **Split Dynamic Disk Group** from the context menu. You can also select the command from the **Actions** menu on the **VEA console** menu bar.
- 3 The **Dynamic Disk Group Split** wizard appears. Click **Next** to continue.
- 4 In the screen that appears next, specify the information that is needed for splitting the disk group as follows:
 - Enter the new dynamic disk group name.

Note: The disk group name should be unique. The dynamic disk group name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

- Choose either the "Split dynamic disk group by disks" option or the "Split dynamic disk group by volumes" option.

When you choose to split by disks, disks appear in the panes in the lower section of the window. If you choose to split by volumes, volumes appear instead of disks.

- Select the disks or volumes that you want to split.
The currently available disks or volumes in the specified disk group are displayed in the left pane, which is labeled "Available disks" or "Available volumes." You select disks or volumes in the left pane and then click **Add** to move them to the right pane, which is labeled "Selected disks" or "Selected volumes.

"You can use **Add All** to move more than one disk at once. If all disks are moved, you need to move at least one of them back to the left pane, because you cannot split a dynamic disk group if you include all the disks in the target disk group.

The volumes and disks that belong to the source dynamic disk group and that are not split remain accessible during the entire dynamic disk group split operation.

After entering the specifications, click **Next** to continue.

5 Verify the information for the dynamic disk group split.

The screen that follows shows the type of information you may see. The purpose of this screen is to let you confirm the disk group split.

The wizard shows the disks that are selected for the split in the left column and the volumes that are split in the right column. In this screen, you may see more disks than what you selected for the split. If the disks you have selected result in a disk group with partial volumes, the program supplies the disks needed to ensure that the split-off dynamic disk group have complete volumes.

If you want to make the split-off dynamic disk group larger, you can click **Back** to go back and add more disks.

When you are satisfied that the information is correct, click **Next**.

6 Click **Finish** to confirm the dynamic disk group split.

If the dynamic disk group split is successful, you'll be able to view the new target dynamic disk group in the tree view and in the right pane of the console. By default, the new target disk group is in the Imported state if you use the GUI to perform the split. If you use the command line to execute the split, the new target disk group is in the Deported state by default, because it assumes that you want to deport the disk group and then import it on another computer. However, with the command line, you can use the `-i` switch with the `vx dg split` command to have the disk group remain in the Imported state.

If the Split Dynamic Disk Group command fails, an error dialog box is displayed showing the reason for failure. The dynamic disk group split operation fails if the target disk group already exists or if a problem occurs when the split operation takes place.

If the computer fails during the split operation, Storage Foundation provides a recovery process. The next section has the details on the recovery process.

Recovery for the split command

If disk and, or system failures occur during the disk group split operation, Storage Foundation generally performs any necessary recovery operations automatically. The recovery procedure performs any operations necessary to restore the disk group to a consistent state. On successful recovery, the disk groups are again available for use.

For example, if there is a power outage during the split operation, it results in incomplete information on the disk and in the disk group. The recovery mechanism ensures that when the host boots up the next time or if the disk group is imported on a host (the same host or another host), the inconsistencies in the information are removed. Depending on the amount of processing completed by the split

operation at the time of the failure, the recovered disk group is either in the state it was in before the split operation or be successfully split into two disk groups.

Recovery has the following two modes:

- Automatic recovery

If it is possible to restart the system, Storage Foundation attempts automatic recovery after the restart. If the automatic recovery process is able to detect sufficient information to perform the recovery, the recovery is successful and the disk group appears in a normal state after the automatic recovery process.

- Manual recovery

In the following situations, it is necessary to do manual recovery:

- If the automatic recovery runs after the restart and there is not sufficient information for recovery, the disk group is in a Deported state and displays an alert icon (a yellow caution icon).
- If you cannot restart the system because it has failed completely, then you must move the disk group's disks to another system and then do a manual recovery. In this situation, it is recommended that you move the entire source disk group to another computer, rescan, and implement the recovery manually.
- It becomes clear that a disk group needs manual recovery after a split command when Storage Foundation does not allow imports of the disk group to succeed or it blocks the join of a disk group. You receive an error message indicating that the disk group needs recovery.

Recovery should be done as soon as possible to preserve the integrity of the data.

You can perform manual recovery through the GUI as described in the following procedure.

You can also perform manual recovery through the command line.

See [“vx dg recover”](#) on page 615.

On the successful completion of the manual recovery operation, you can perform all Storage Foundation operations. The only exception is that if a host crashes during a particular stage of the split operation, one or more of the disk group's disks and volumes may come up in an unhealthy state after recovery. In that situation, use the Reactivate Disk command on the unhealthy disks and then the Reactivate Volume command on each of the unhealthy volumes that are involved in the split command. After running these commands, the disk group's disks and volumes should be healthy.

See [“Reactivate disk command”](#) on page 540.

See [“Reactivate volume command”](#) on page 541.

To manually recover the dynamic disk group

- 1 Select the disk group in the VEA GUI and select **Recover Dynamic Disk Group** from the context menu.
- 2 In the dialog box that appears, make sure that the proper disk group name is displayed, and click **OK** to have Storage Foundation start the recovery process.

Dynamic disk group join

This section describes the procedures for a dynamic disk group join operation.

When you join two dynamic disk groups, the disk group you designate as the source becomes the one that loses its identity in the merge. The target dynamic disk group is the one that remains after the merge.

Storage Foundation allows a partial join—that is, volumes in the source disk group can have missing or nonfunctional disks and the join command can still succeed.

See [“About partial disk group join”](#) on page 281.

The disk group type after the join becomes the type of the target disk group. For example, if the target disk group before the join had private dynamic disk group protection, the combined disk group has private dynamic disk group protection after the join.

Note: A cluster dynamic disk group that is part of the cluster resources cannot be a source disk group for a join command. However, it can be a target disk group for the command.

Note: It is possible to force the join command when a volume is in use or if disks from the source disk group’s disks are missing, but this is not recommended as a best practice.

Note: Microsoft Disk Management Disk Groups do not support dynamic disk group join.

To join two disk groups into one disk group:

- 1** Make sure that the volumes in the source dynamic disk group are not in use and are healthy.
- 2** Right-click the disk group object in the left pane of the VEA console, and then click the **Join Dynamic Disk Group** command, or select the **Join Dynamic Disk Group** command from the **Actions** menu on the **VEA console** menu bar.

- 3 In the **Join Dynamic Disk Group** dialog box, make sure that the correct source and target disk groups are indicated. If necessary, use the drop-down lists to change the disk group names.

The **source dynamic disk group** is the disk group that is merged with the target dynamic disk group. The source disk group ceases to exist after the join.

The **target dynamic disk group** is the disk group that remains after the join. It retains its identity and becomes larger with the addition of the disks from the source disk group.

The **Clear host ID** check box should be used when you want to import a disk group from another node and that disk group was not properly deported. A host ID might not be cleared, for example, if the host computer became nonfunctional. When a host ID has not been cleared, SFW does not permit the join without the Clear host ID override. Use this override with care.

The **Allow Partial Join** check box can be used when you have a missing or nonfunctional disk or disks in the source disk group and you want SFW to allow the join to occur.

See See [“About partial disk group join”](#) on page 281.

- 4 Click **OK**.

If the join operation is successful, the source dynamic disk group merges into the target dynamic disk group. The changes in the target dynamic disk group are reflected in the VEA console tree view and the right pane.

If a volume is in use, SFW displays an error message about forcing the join command.

See [“Error message about forcing a split or join operation”](#) on page 286.

If there is an error message about the disks not being ready, you can click the check box for the partial join to allow the command to complete.

If the join operation fails, an error dialog box is displayed, showing the reason for the failure.

The new dynamic disk group after the join command are of the same type as the target dynamic disk group. For example, if a cluster dynamic disk group is joined with a normal dynamic disk group, then the new dynamic disk group is a normal dynamic disk group.

About partial disk group join

Storage Foundation facilitates for a disk partial join—that is, volumes in the source disk group can have missing or nonfunctional disks and the join command still succeeds. However, you need to click the **Allow Partial Join** check box in the **Join**

Dynamic Disk Group window for the join to take place. When the missing disk(s) come back to an Online state, then you need to do another Join Dynamic Disk Group command to add the missing disk or disks to the target disk group.

Warning: If you have a partial join in place, you must wait until the missing disks or disks are joined to the target disk group before making the following configuration changes to either the source or target disk group: replace or evacuate disk, split or move a subdisk, extend a volume, add a mirror, and clear hot relocation. Doing so before all disks are joined can result in data loss.

Using Dynamic Disk Group Split and Join with a cluster on shared storage

This section describes the use of the Dynamic Disk Group Split and Join operations with SFW cluster disk groups on shared storage. SFW supports cluster dynamic disk groups with either Cluster Server (VCS) or Microsoft Failover Clustering.

See [“Using Dynamic Disk Group Split and Join with VCS or Microsoft Failover Clustering”](#) on page 282.

See [“Example of an off-host backup procedure using SFW and VCS or Microsoft Failover Clustering”](#) on page 284.

Using Dynamic Disk Group Split and Join with VCS or Microsoft Failover Clustering

If a disk group is under Cluster Server (VCS) or Microsoft Failover Clustering control, SFW imposes certain conditions on the functioning of the Dynamic Disk Group Split and Join operations.

These conditions are the following:

- If a shared dynamic cluster disk group is joined to another (possibly shared) dynamic disk group and the source disk group is configured as a VCS or Microsoft Failover Clustering resource, then the join operation fails.
- If a dynamic disk group (possibly shared) is joined to another shared cluster dynamic disk group and the target disk group is configured as a VCS or Microsoft Failover Clustering resource, then the join operation is allowed to proceed.
- (VCS only) If a shared dynamic cluster disk group is split into another (possibly shared) disk group and the source disk group is configured as a VCS resource, then the split operation is not allowed to proceed, with one exception. The exception is that the split operation can proceed when the volumes selected to be split off are not configured under VCS.

Thus, in this situation, if users want to split any volumes that are configured under VCS, they should remove the volumes from VCS monitoring and retry the split operation.

- (Microsoft Failover Clustering only) If a shared dynamic cluster disk group is split into another (possibly shared) disk group and the source disk group is configured as a Microsoft Failover Clustering resource, the split operation is allowed to proceed. However, before the command completes, a confirmation window appears to remind you that if you split a disk group resource so that a volume on which other cluster resources depend is moved to a new SFW cluster disk group, it is necessary to define the new cluster disk group as a Microsoft Failover Clustering resource and to modify the properties of dependent resources to reflect the change.

If necessary, use Windows Server's Failover Cluster Manager to check and modify resource dependency fields as follows:

- Define the new cluster disk group as a Volume Manager cluster disk group resource.
Even though the name of the program has changed to Storage Foundation, the resource group is still named "Volume Manager."
- If the new disk group resource is in a different Microsoft Failover Cluster group than the original cluster disk group resource, move all resources that depend on the new cluster disk group resource to the Microsoft Failover Cluster group where it resides.
- For each dependent resource, do the following:
 - Offline the resource.
 - Select the resource and select **Properties**.
 - Select Dependencies in the resource's Properties window.
 - If the resource lists the disk group that was the source of the split operation as a dependency, click **Modify**.
 - Move the new disk group from the **Available Resources** pane of the Modify Dependencies window to the **Dependencies** pane, and move the old disk group from the **Dependencies** pane to the **Available Resources** pane, if it appears in the **Dependencies** pane.
 - Online the resource.
- (VCS only) If there is a hardware failure or a power failure during an ongoing disk group split operation and the disk group being split is configured under VCS, then VCS attempts to failover that disk group to another node in the cluster and brings the disk group online.

At that time, appropriate disk group split recovery is carried out as required to bring the disk group online.

See [“Recovery for the split command”](#) on page 277.

Manual intervention may or may not be required to recover the disk group.

- If you are splitting a dynamic cluster disk group that contains the quorum volume, make sure that the quorum volume is not part of the target disk group after the split operation. If you fail to do so, the quorum volume is not able to function and the cluster fails.

Example of an off-host backup procedure using SFW and VCS or Microsoft Failover Clustering

This section describes how to use SFW and VCS or Microsoft Failover Clustering in a typical off-host backup procedure.

In the example, a disk group named "SQL" has volume "V" on which a SQL Server database is under heavy use by clients. This disk group is imported on one node of a cluster running VCS or Microsoft Failover Clustering. VCS or Microsoft Failover Clustering monitors the disk group as well as the volume, ensuring high availability.

In the off-host backup solution that is outlined below, the high availability of the volume V is not affected at all. While the successful backup takes place, the SQL server provides uninterrupted service to its clients with minimum configuration changes on both nodes in the cluster.

The procedure is fully scriptable and thus can be automated to provide a backup on an hourly, daily, weekly, or monthly basis.

To implement off-host backup

- 1 Create a snapshot "V1" of the volume V with the snap commands in the VEA GUI or the CLI.

See [“Snapshot commands”](#) on page 259.
- 2 Split the dynamic disk group SQL into another dynamic disk group "Temp" so that V1 resides on the disks in the Temp disk group. Since VCS or Microsoft Failover Clustering monitors V (and not V1), the split operation succeeds.

See [“Dynamic disk group split”](#) on page 274.
- 3 Deport the Temp disk group and import it on another node in the cluster.

See [“Importing and deporting dynamic disk groups”](#) on page 157.
- 4 Back up volume V1 on the second node, using the backup software.
- 5 After backup is complete, deport the Temp disk group on the second node.

- 6 Join the Temp disk group to the SQL disk group on the first node. Since the SQL disk group is a VCS or Microsoft Failover Clustering resource (and the Temp disk group is not monitored), the Join operation succeeds.
See [“Dynamic disk group join”](#) on page 279.
- 7 The snapshot volume V1 can then be snapped back to its parent volume V.
See [“Snap back”](#) on page 265.

Limitations when using dynamic disk group split and join with Volume Replicator

If you use the Dynamic Disk Group Split and Join commands with the volumes that are under Volume Replicator (Volume Replicator) control, some limitations apply.

These include the following:

- A volume under replication cannot be part of the volumes that are split off to the target disk group. It can remain in the source disk group.
- A join operation is not allowed if the source disk group has any Volume Replicator objects.

Dynamic Disk Group Split and Join troubleshooting tips

This section provides information and workarounds for the problems that might arise when you use the Dynamic Disk Group Split and Join commands.

Error message, “The disk group requires recovery. please recover...”

This message can occur when you attempt to import a disk group or do a join command. The message indicates that the disk group needs recovery because of a problem that occurred during a disk group split operation.

See [“Recovery for the split command”](#) on page 277.

Error message, “One or more disk(s) in the source dynamic disk group are missing.”

This message occurs when you attempt a disk group split or join operation and a disk in the disk group is no longer functional or has gone offline. It is recommended that you rescan and then retry the disk operation. However, if one or more disks in the disk group are missing or nonfunctioning, you can click the **Allow Partial Join** check box to have the join command succeed even though one or more disks are missing.

See [“About partial disk group join”](#) on page 281.

Error message about forcing a split or join operation

If you attempt a dynamic disk group split or join command when a volume is in use, Storage Foundation puts up a message asking whether you want to force the split or join. If the volume is in use because it is selected in Windows Explorer, then you can force the split or join safely. However, if the volume is in use within an application, it is not recommended that you attempt to complete the command. If a volume is regenerating, resynching, or recovering, then even if you force the split or join, the operation is not successful.

CLI FlashSnap commands

SFW provides CLI commands corresponding to the VEA GUI FlashSnap commands. The following table lists each GUI FlashSnap command with a cross reference to its corresponding CLI command. It also has some additional comments to clarify distinctions between the two CLI snapshot commands, `vxsnap` and `vxassist snapshot`.

The following are the FlashSnap commands:

Table 11-1 FlashSnap GUI and CLI Commands Correspondences

GUI Command	Corresponding CLI Command
FastResync	See “vxvol set fastresync=on off” on page 583.
Prepare	See “vxassist prepare” on page 664. Note: The Prepare command replaces the Snap Start command in the GUI. Both <code>prepare</code> and <code>start</code> keywords are available in the CLI, however <code>prepare</code> is the recommended keyword.
Snap Shot	There are two types of snapshot commands. <ul style="list-style-type: none"> ■ The command <code>vxassist snapshot</code> allows snapshotting only a single volume at a time. See “vxassist” on page 638. ■ The command <code>vxsnap</code> allows simultaneous, multiple snapshots. See “vxsnap” on page 710. The <code>vxsnap</code> command integrates with VSS to enable VSS snapshots.
Snap Back	See “vxassist snapback” on page 669.
Snap Clear	See “vxassist snapclear” on page 670.

Table 11-1 FlashSnap GUI and CLI Commands Correspondences
(continued)

GUI Command	Corresponding CLI Command
Snap Abort	See “vxassist snapabort” on page 672.

Note: For information on how to use FlashSnap to perform snapshots and implement recovery procedures, refer to the SFW HA Quick Recovery (QR) Solutions guides.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist prepare`, `vxassist snapshot`, `vxassist snapback`, `vxassist snapclear`, and `vxassist snapabort` commands.

Fast File Resync

Fast File Resync (FFR) lets you resynchronize one or more files in a snapshotted volume back to the original volume.

This section covers the following topics:

See [“Overview”](#) on page 287.

See [“Prerequisites”](#) on page 288.

See [“Resynchronizing a file”](#) on page 288.

See [“Possible error for small files”](#) on page 289.

Overview

FFR uses Veritas FastResync technology to resynchronize selected files in a snapshotted volume to the original volume. Fast File Resync is designed for database applications using volumes with multiple database files and can be used to restore one or more corrupt files. A snapback command that resynchronizes the split-mirror snapshot back to the original files is performed for the selected files. Other databases do not have to be off-lined during the Fast File Resync process, and the entire volume is not resynchronized. The Fast File Resync command restores the file or files on the original volume but does not bring the database to a current state.

In general, the FastResync technology is more efficient for files greater than 50 MB, and the Windows Explorer copy function or Windows `xcopy` command is more efficient for files smaller than approximately 50 MB. FFR is unable to restore a file that has been deleted.

Note: FFR does not work on files smaller than 8 KB.

See [“Possible error for small files”](#) on page 289.

Warning: It is important to wait until the FFR process is complete before you access and use the restored file. Data corruption can occur if the file is used before the resynchronization is complete.

The Fast File Resync command can be issued through the VEA or through a CLI command, `vxfsync`.

However, the CLI command does not support resynchronization of multiple files. It supports the resynchronization of one specified file at a time.

See [“vxfsync”](#) on page 743.

Note: FFR is not available through the VEA GUI of a SFW remote client that is connected to a SFW server.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Fast File Resync.

Prerequisites

Prerequisites must be met for a successful FFR resynchronization of a file.

These include the following:

- The names and creation times of the snapshot and original files are identical.
- The file is formatted using NTFS.
- The file is not compressed.
- The file is not encrypted.
- The files do not have any open handles.

Resynchronizing a file

The following is the procedure to resynchronize a file.

To resynchronize corrupted files

- 1 Right-click on the original volume containing the corrupted file and select **Fast File Resync**.

The Fast File Resync dialog box appears.

- 2 Click **Browse** and navigate to the file or files that you want to copy from the snapshot volume back to the original volume. Use the Control and Shift keys to select multiple files. Click **Select**.

- 3 Select the desired snapshot volume from the **Snap list**.

A snapshot volume must have either a drive letter or a mount point assigned for it to be available for selection from the **Snap list** in the **Fast File Resync** window.

- 4 Click **OK**.

Possible error for small files

When using FFR to resynchronize files smaller than 8 KB, the file extents call to the operating system may fail and an error message notifies you that the call has failed and refer you to the trace for further information.

If this error occurs, use the Windows Explorer copy function or Windows xcopy command to copy the file.

The trace can be found at `\Veritas\Veritas Object Bus\logs\vxisis.log`.
The error is prefixed by `vxffr`.

Volume Shadow Copy Service (VSS)

This section describes how Storage Foundation works with Volume Shadow Copy Service (VSS) to provide snapshots.

This section covers the following topics:

See [“Overview”](#) on page 289.

See [“VSS components”](#) on page 290.

See [“Overview of VSS process for SFW snapshots”](#) on page 292.

See [“VSS snapshot wizard for Microsoft Exchange and Microsoft SQL”](#) on page 293.

Overview

Volume Shadow Copy Service (VSS) is a Windows service that provides the capability of creating snapshots or volume shadow copies. A volume shadow copy

is a volume that represents a duplicate of the state of the original volume at the time the copy began. Thus, through VSS, point-in-time copies of data can be saved and used for different purposes, such as backup or application testing.

VSS snapshots represent an advance in snapshot technology because snapshots can be taken when application files are open. VSS interacts with an application (such as Microsoft SQL Server) and can quiesce the application for the moment when the snapshot is created. VSS restarts the application immediately after the snapshot. VSS only takes snapshots of read or write volumes and the resulting volume shadow copy is a read-only volume.

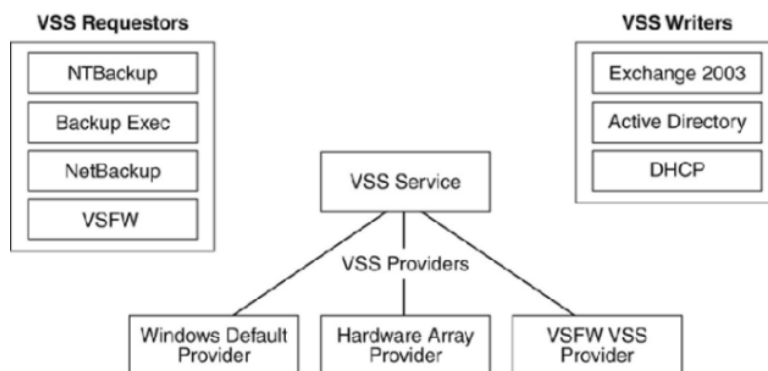
Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Volume Shadow Copy Service (VSS).

VSS components

VSS works through an infrastructure that coordinates the snapshot process by using software components from Microsoft itself and from various storage software and hardware vendors.

The following figure indicates the VSS service and the three key software components that are used in the VSS process—requestors, writers, and providers.

Figure 11-1 VSS components



Requestors

The software component that requests the creation of the snapshot is known as the requestor. NetBackup is the default requestor that comes with Windows. Backup Exec and NetBackup also have requestor components for VSS. In addition, SFW also functions as a VSS requestor.

Writers

These are application-specific software components that ensure that the application data is quiesced and then restarted after the snapshot. Microsoft has provided the writer software component in SQL Server and in the Active Directory and DHCP software features that are included in Windows.

Providers

Providers are the software components that implement the snapshots. Windows comes with a default provider. However, the SFW VSS provider and providers from hardware array manufacturers offer enhanced functionality.

For the VSS process to work properly, the different software programs that are involved in the process need to be "VSS-aware" — that is, the necessary software components have been added to the programs so that they can participate. Again, it may take time for software and hardware vendors to provide the various VSS components with their products.

SFW as a VSS Provider and Requestor

An example in which SFW acts as a provider is a backup situation where Backup Exec is the requestor, initiating a snapshot of a SFW volume that is part of a Microsoft Exchange 2010 mailbox database. Microsoft Exchange is the VSS-aware writer.

As a VSS requestor, SFW with FlashSnap can initiate snapshots at any time. It is not dependent on a backup program to initiate the snapshot. Thus, SFW with FlashSnap integrates with VSS and is able to snapshot the volumes that are associated with an Exchange 2010 mailbox databases without taking the databases offline. The VSS process allows the databases of the storage group to be quiesced before the snapshot operation occurs and then reactivated immediately after it. This quiescing, supported by Exchange 2010 at the mailbox database level, allows for Microsoft supported and guaranteed persistent snapshots of your data.

Recovery

Through SFW's VSS Snapshot wizard or `vxsnap restore` command, the VSS snapshots can be used for a point-in-time recovery of databases or a roll-forward recovery to the point of failure of either the mailbox databases or individual database within it.

In Storage Foundation, these different types are supported using the VSS Restore operation.

The point-in-time recovery restores the data to the point in time that the quick recovery image was last refreshed—that is, it restores all the data to what was recorded in the snapshot.

The point-of-failure recovery recovers the SQL Server databases by restoring the old image of the data volumes only and replaying the logs to accomplish a roll-forward recovery, provided that the log volume is intact.

Refer to the application-specific Quick Recovery Solutions guide for detailed procedures on how to use FlashSnap with the application, such as Microsoft SQL Server, to perform snapshots and to implement recovery procedures.

See “[vxsnap](#)” on page 710.

Advantage of SFW Snapshots

The snapshots that are created with VSS through the Windows default provider are limited to copy-on-write snapshots. That is, the snapshot does not include a full copy of the original image but refers back to the original volume. For this reason, the default provider snapshots cannot be used for any procedures that involve off-host processing. The Veritas FlashSnap snapshot procedures produce independent split-mirror snapshot volumes that can be used for off-host processing. Thus, the FlashSnap snapshots that are implemented through SFW offer a definite advantage over the snapshots that are created with the Windows default provider.

Overview of VSS process for SFW snapshots

The process for creating SFW snapshots with VSS can be done through the GUI, the command line, or a script with CLI commands.

Note: Snapshots can only be taken of read or write volumes when VSS is used. The resulting VSS snapshot is read only. The file system label of VSS snapshot volumes cannot be changed.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot or VSS-related commands.

It is recommended that you use the VEAG GUI Snap Back operation or `vxassist snapback` to reattach the snapshot mirror to its original volume. This enables you to update the snapshot mirror and make it available for more timely backups.

To create SFW snapshots with VSS

- 1 Use the **Prepare** command through the VEA GUI to attach a snapshot mirror to a volume on the working server.

The `vxassist prepare` command can also be used.

Creating the snapshot mirror of the volume and resynchronizing it to the volume takes considerable time, but it only has to be done once for the initial setup.

The **Prepare** command replaces the **Snap Start** command in the GUI.

Both `prepare` and `start` keywords are available in the CLI, however `prepare` is the recommended keyword.

- 2 Request a snapshot through a backup program (such as Backup Exec) or through SFW. In SFW, initiate a VSS snapshot through the GUI with the **Snap Shot** command. Be sure to enable the VSS snapshot method in the Systems Settings.

You can also use the `vxsnap` CLI command to take the snapshot.

If you plan to use these snapshots to recover a VSS-aware application's database, use the VSS Snapshot wizard.

You can also use the `vxsnap` command.

See [“Using the VSS snapshot wizards with Microsoft Exchange”](#) on page 295.

The `vxsnap` command lets you snapshot multiple volumes simultaneously; thus, you can snapshot all the volumes in a VSS-aware application's storage group at one time. The VSS Snapshot wizard can also snapshot a storage group one at a time.

- 3 VSS interacts with a writer utility for the application to make sure that the application is momentarily quiesced for the snapshot.
- 4 VSS determines the appropriate provider for the snapshot. The Microsoft default provider is used only if another provider is not present. If an SFW dynamic volume is involved, then SFW is selected as the provider.
- 5 After the snapshot is taken, the writer utility makes sure that the application is restarted.
- 6 Once the snapshot takes place successfully, VSS communicates the successful result to the program with the requestor software.

VSS snapshot wizard for Microsoft Exchange and Microsoft SQL

SFW integrates with the Windows Volume Shadow Copy Service (VSS) as both a VSS Requestor and a VSS Provider. FlashSnap provides this integration.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot or VSS-related commands.

FlashSnap integrates with VSS to provide the ability to detach multiple split-mirror snapshot volumes simultaneously. This lets you snapshot all the volumes that are associated with an Exchange storage group or SQL database at exactly the same point in time without taking the databases offline. When viewed all together these snapshots form a snapshot set of the storage group, component, or database. These persistent FlashSnap snapshots can later be used for either a point-in-time recovery or a roll forward recovery to point of failure of the storage group or database.

FlashSnap calls VSS to perform a "copy backup" of Exchange, which does not truncate the transaction log files.

The following occur during the snapshot process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a split-mirror snapshot of an Exchange storage group or SQL database.
- The VSS coordinator service calls the Exchange VSS Writer or SQL VSS Writer to find out which volumes contain the databases and transaction logs.
- The VSS coordinator service notifies the FlashSnap VSS Provider to prepare for the snapshot.
- Acting as a VSS Requestor, FlashSnap requests that the VSS coordinator service begin the snapshot call.
- The VSS coordinator service notifies the Exchange VSS Writer or SQL VSS Writer to quiesce the databases in preparation for the snapshot. When this is accomplished, the Exchange Writer or SQL Writer informs the VSS coordinator service to proceed.
- The VSS coordinator service calls the FlashSnap Provider to create the split-mirror snapshot by detaching the snapshot volume from the original volume. After the snapshot volume is detached, the FlashSnap Provider informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the Exchange Writer or SQL Writer to resume normal I/O.

Using the VSS snapshot wizards with Microsoft Exchange

SFW provides support for taking snapshots of Exchange 2007 storage groups and Exchange 2010 mailbox databases. FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all the volumes that are associated with an Exchange mailbox database without taking the databases offline. The VSS Snapshot wizard uses VSS to take snapshots of mailbox databases. The VSS process allows the databases to be quiesced before the snapshot operation and reactivates them immediately after. The quiescing of the storage groups (Exchange 2007) mailbox databases (Exchange 2010) level allows for Microsoft supported and guaranteed persistent snapshots of the data. A snapshot of a storage group or mailbox database can be reattached and resynchronized to match the current state of the storage group or mailbox database respectively with the VSS Snapback wizard.

Note: Snapshots can only be taken of read or write volumes when VSS is used. The resulting VSS snapshot is read only. The file system label of VSS snapshot volumes cannot be changed.

SFW also provides a VSS Snapshot Scheduler wizard that can be used as an alternative to the VSS Snapshot wizard and the VSS Snapback wizard. It enables you to set up a schedule for taking the initial snapshots and for automating the snapback process. At the scheduled time for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again. The schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

SFW also provides recovery support for an Exchange storage group or a selected database within a storage group for Exchange 2007 and also for the Exchange 2010 mailbox database. Using the VSS Restore Wizard, the snapshots that are taken with the VSS Snapshot wizard can be used for a point-in-time recovery of the storage group or a roll-forward recovery to the point of failure of either the storage group or an individual database in the storage group.

In case of Exchange 2010, mailbox databases are no longer tied to a particular server. Database configuration information is stored at the same hierarchy level as the Exchange configuration objects. Database names have to be unique across the organization. The failover occurs at the database level not the server level as in previous versions of Exchange.

The point-of-failure recovery recovers the Exchange storage group or mailbox database by restoring the old image of the data volumes only and replaying the logs to accomplish a roll-forward recovery, provided that the log volume is intact.

The point-in-time recovery restores the data to the time that the snapshot set was last refreshed—that is, it restores all the data that was recorded in the snapshot.

For Exchange 2007 only, if you have set up a recovery storage group (RSG), you have the option to restore to the RSG rather than to the production volume, leaving the production volume intact. Refer to Microsoft instructions and documentation in creating an RSG. The database names in the recovery storage group must match those in the source storage group.

Refer to the *Storage Foundation and High Availability Solutions High Availability and Disaster Recovery Solutions Guide for Microsoft Exchange* for additional information about how to use FlashSnap with Microsoft Exchange to perform and to implement recovery procedures.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot or VSS-related commands.

See [“Using the VSS Snapshot wizard”](#) on page 296.

See [“Using the VSS Snapback wizard”](#) on page 300.

See [“Using the VSS Snapshot Scheduler wizard”](#) on page 301.

See [“Using the VSS Restore Wizard ”](#) on page 309.

See [“Refreshing the snapshot set”](#) on page 319.

Using the VSS Snapshot wizard

Creating a snapshot is a two-step process. The first step, prepare the volume for snapshot, creates snapshot mirrors attached to all the original volumes in the specified storage group. Depending on the size and number of volumes, the process of synchronizing the new snapshot mirrors with the original production volumes can take a long time. The second step uses the VSS Snapshot wizard to create the snapshot set by detaching the snapshot mirrors from the original volumes and creating separate on-host snapshot volumes as well as an XML file to store the Exchange and snapshot volume metadata. The VSS Snapshot wizard integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes in the storage group. This snapshot is done while the databases are online and without disrupting the email flow. Once a snapshot set has been created, it can be reattached and resynchronized with the VSS Snapback wizard.

Note: After performing operations for snapshots, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

Creating the snapshot set

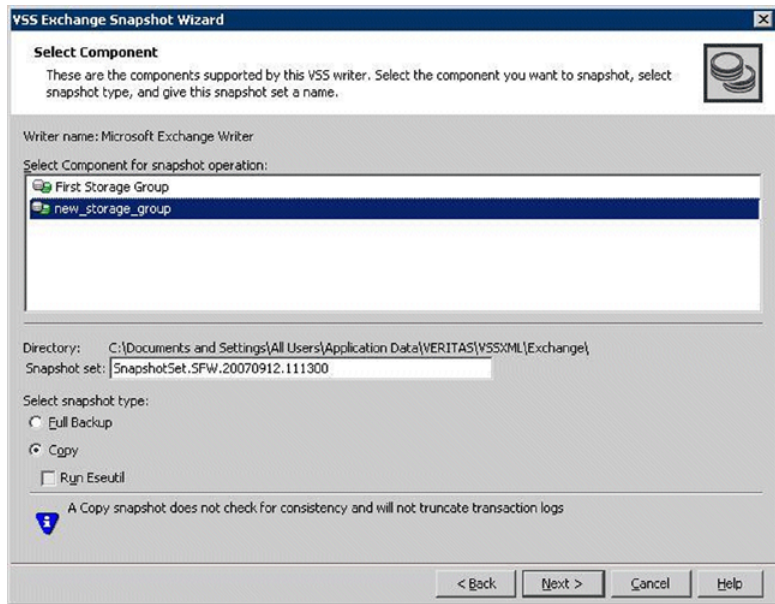
The following is the procedure to create a snapshot set.

Note: You must prepare each volume for a snapshot before running the VSS Snapshot wizard.

To create the snapshot set using the VEA GUI

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Right-click **Exchange** and click **VSS Snapshot**.
 - (Exchange 2007 only) Optionally, if replication is enabled and you want to create a snapshot of the replica, right-click **Exchange Replica** and click **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a snapshot on the replication node, right-click **Exchange Replica** and click **VSS Exchange Snapshot**.

- 4 In the wizard, review the **Welcome** page and click **Next**.



- 5 Specify the snapshot set parameters as follows and then click **Next**:

Select Component for snapshot operation Select the component for the snapshot set.

- For Exchange 2007, the component is a storage group.
- For Exchange 2010, the component is a mailbox database.

Using the VSS snapshot wizards with Microsoft Exchange

Snapshot set	<p>Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.</p>
	<p>The XML file is stored by default in the directory that is shown on the screen. However, in a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, use a text editor to create a text file named "redirect.txt." This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the metadata file, for example, \\ServerName\SharedFolderName. Save the redirect.txt file in the default directory C:\Program Files\Veritas\Veritas Volume Manager\VSSXML on each node of the cluster.</p> <p>Note: You must not use the volume name or volume path in the redirect.txt file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore fails.</p> <p>Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over. You can also specify the file share path in the redirect.txt file.</p>
Select snapshot type	<p>Select the snapshot type.</p> <p>Full Backup is typically used for backup to tape or other storage media. It does the following:</p> <ul style="list-style-type: none">■ Creates a copy of the selected component■ Runs Eseutil to check for consistency before truncating the logs■ Truncates the transaction logs <p>Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally check Run Eseutil with the Copy option to check the snapshot for consistency.</p>

- 6 In the **Change Attributes** panel, optionally change the attributes for the snapshot volumes as follows and click Next:

Snapshot Volume Label	Displays the read-only label for the snapshot volume.
Drive Letter	Optionally, click a drive letter and select a new choice from the drop-down menu.
Plex	Optionally, click a plex and select a new choice from the drop-down menu.

- 7 Optionally, in the **Synchronized Snapshot** panel, select the secondary hosts for which you want to create synchronized snapshots.

Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Volume Replicator (Volume Replicator). Otherwise, you are directly taken to the **Schedule Information** panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

Using the VSS Snapback wizard

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot set so that it matches the current state of its original Exchange database. The wizard is available in the context menu of the VSS Writer object.

Resynchronizing the snapshot set

The following is the procedure to resynchronize (snapback) a snapshot set.

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the snapshot set.
- 2 From the VEA, select the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click the node of the selected application and click **VSS Snapback**.

- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.

The XML metadata file contains all the information that is required to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.
- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click **Yes** to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

Using the VSS Snapshot Scheduler wizard

You can use the VSS Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of refreshing. At the scheduled time for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again. Once configured and applied, the schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background. In a clustered server environment, ensure that the scheduler service is configured on each node with domain administrator privileges so that any node in the cluster can run the schedule after a failover.

In case of a VCS cluster setup with more than one node, configure the Veritas Scheduler Service with any user account (other than the Local System account) which is valid on all the nodes of the cluster. This user account which is configured for the Scheduler Service must have read-write permissions to the file share path.

Note: Information about the operations that the schedule performs can be found by expanding the nodes under the Scheduled Tasks node in the tree-view of the VEA GUI. Expand the appropriate application node to locate the schedules for that application. Right-click the node and select **Job History** to view information about the operations performed by the schedule.

The VSS Snapshot Scheduler wizard does not prepare the snapshot mirror. Prepare the snapshot mirror with the **Prepare** command before running the VSS Snapshot Scheduler wizard.

Note: The Prepare command replaces the Snap Start command in the VEA GUI.

Note that if Exchange 2010 mailbox database is configured under VCS, then you can store the snapshot set XML metadata file in a file share path by configuring a file share resource. This is to avoid configuring extra shared volumes to store the snapshot set file, which is available once the mailbox database fails over.

There can be large number of mailbox databases in Exchange 2010 and each database may have independent schedules. Hence the number of different paths that are required for storing the snapshot metadata files are higher, which can be achieved by configuring a file share resource.

In case of a cluster setup, even though schedules are visible on all nodes, snapshot happens only on a node where mailbox database is present.

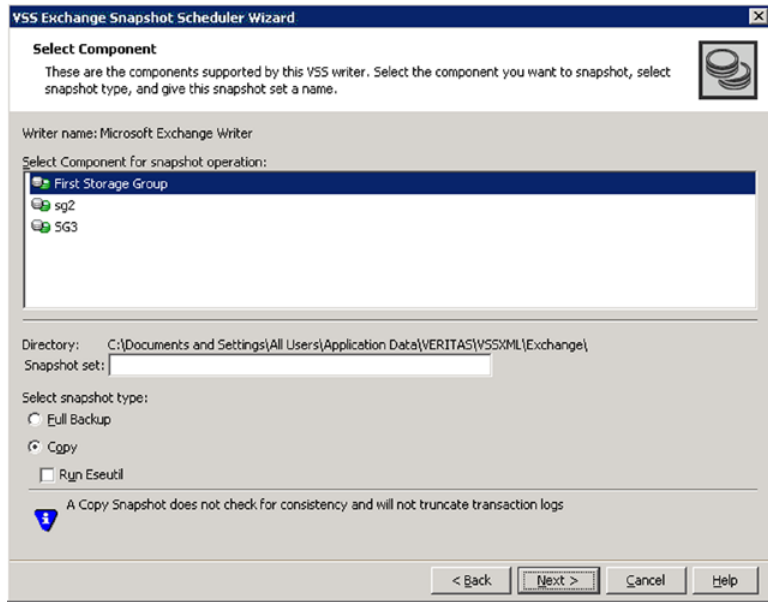
Note: For Exchange 2010, before adding or removing a node from a VCS cluster setup, delete the schedules and then recreate the schedules on the required node.

Note: For an Exchange 2010 VCS cluster setup if a schedule is created on a file share, then it is visible on all nodes of the cluster and can be deleted from any node irrespective of where the Exchange mailbox database component is online.

To schedule a snapshot

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Choose one of the following:
 - Right-click **Exchange** and select **VSS Exchange Snapshot**.
 - Optionally, if replication is enabled in an Exchange 2010 DAG and you want to create a schedule on the replication node, right-click **Exchange Replica** and click **Schedule VSS Exchange Snapshot**.

4 In the **Welcome** panel, review the information and click **Next**.



5 Specify the snapshot set parameters as follows and then click **Next**:

- | | |
|---|--|
| Select component
for snapshot
operation | Select the component for the snapshot set.

For Exchange 2007, the component is a storage group.

For Exchange 2010, the component is a mailbox database. For
Exchange 2010, the component is a mailbox database. |
|---|--|

Snapshot set	<p>Accept or enter a directory location for the XML file. The XML file is stored by default in the directory that is shown on the screen.</p> <p>Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore fails.</p> <p>However, in a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, either edit the directory path in the Directory field for this wizard screen or use a text editor to create a text file named <code>redirect.txt</code>. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, <code>\\ServerName\SharedFolderName</code>. Save the <code>redirect.txt</code> file in the default directory <code>C:\Program Files\Veritas\Veritas Volume Manager\VSSXML</code> on each node of the cluster.</p> <p>You must not use the volume name or volume path in the <code>redirect.txt</code> file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore fails.</p> <p>Note that if your Exchange 2010 mailbox database is configured under a VCS cluster, then you can store the snapshot set metadata file in a file share path by configuring a file share resource. You can also specify the file share path in the <code>redirect.txt</code> path.</p>
Select snapshot type	<p>Select the snapshot type.</p> <p>Full Backup is typically used for backup to tape or other storage media. It does the following:</p> <ul style="list-style-type: none">■ Creates a copy of the selected component■ Runs Eseutil to check for consistency before truncating the logs■ Truncates the transaction logs <p>Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally, check Run Eseutil with the Copy option to check the snapshot for consistency.</p> <p>You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database.</p>

- 6** In the **Change Attributes** panel, optionally change the attributes for the snapshot volumes and click **Next**:

Snapshot Volume Label Displays the read-only label for the snapshot volume.

Drive Letter Optionally, click a drive letter and select a new choice from the drop-down menu.

The specified drive letters may not be available when the snapshot is taken. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

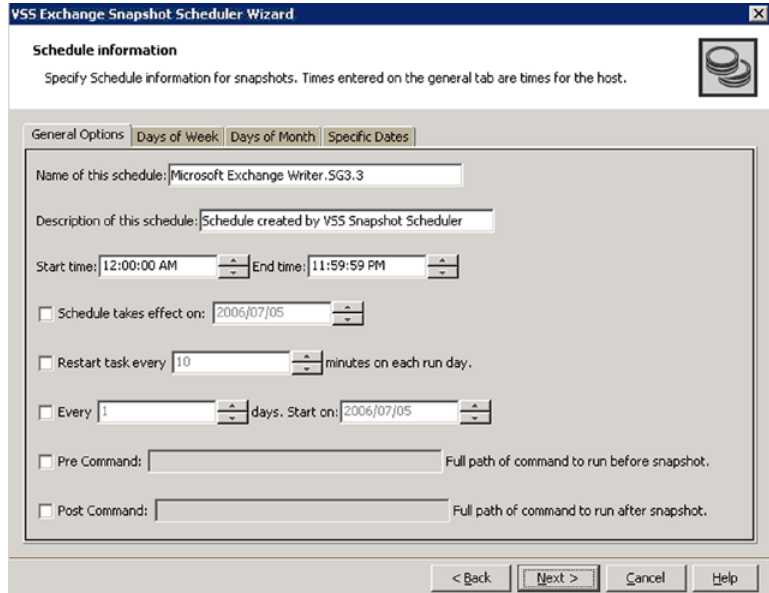
Plex Optionally, click a plex and select a new choice from the drop-down menu.

- 7** Optionally, in the **Synchronized Snapshot** panel (Volume Replicator only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Volume Replicator. Otherwise, you are directly taken to the **Schedule Information** panel.

See *Volume Replicator Administrator's Guide*.

- 8 In the **Schedule Information** panel, on the **General Options** tab, specify the following:



Name of this schedule	Enter a unique name for the snapshot set schedule. This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the VSS writer name and a numbered suffix that increments with each schedule.
Description of this schedule	Optionally, enter a description to help you identify the schedule when you view information about the snapshot status
Start Time	The time of the day to begin taking snapshots.
End Time	The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.
Schedule takes effect on	The date on which the specified schedule takes effect. The default is the current date.
Restart task every	The interval between snapshots, in minutes. For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice.

Using the VSS snapshot wizards with Microsoft Exchange

Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, one day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	If you enable the Every option, specify the starting date.
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p> <p>Note: Pre-command which launch windows or any other GUI-related entities are not allowed in Scheduler.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p> <p>Note: Post-command which launch windows or any other GUI-related entities are not allowed in Scheduler.</p>

9 To specify run days for the schedule, make selections on the following tabs:

Days of Week	<p>Select one or more days on one or more weeks of the month.</p> <p>You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.</p>
Days of Month	<p>Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.</p>
Specific Dates	<p>Select one or more specific dates to include in or to exclude from the schedule.</p> <p>Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.</p>

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday,

- 10 Click **Next**.
- 11 Review the snapshot set and schedule details and click **Finish**.

Displaying the status of the scheduled VSS Exchange snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the **System node** and **StorageAgent node**. Select and click to expand the **Scheduled Tasks > Exchange node**.

The scheduled snapshots are listed on the right pane.

- 3 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Scheduled Tasks** and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.

- 4 In the **Job History** dialog box, view the schedule information.

You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a VSS Exchange snapshot

The following is the procedure to delete a schedule.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the GUI console.

Note: The VEA GUI can delete snapshot schedules only; it does not delete mirror preparation that is scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update the template settings that are created with the Quick Recovery Configuration Wizard.

To delete a schedule from the VEA

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Select and click to expand the **Scheduled Tasks > Exchange node**.
The scheduled snapshots are listed on the right pane.
- 4 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Using the VSS Restore Wizard

SFW provides recovery support for a selected Exchange storage group or a selected database within a storage group for Exchange 2007 storage and Exchange 2010 mailbox database.

Using the VSS Restore Wizard, the snapshots that are created with the VSS Snapshot wizard can be used for a point-in-time recovery of the storage group or a roll-forward recovery to the point of failure of either the storage group or an individual database in the storage group.

The point-in-time recovery restores the data to the time that the snapshot set was last refreshed—that is, it restores all the data that was recorded in the snapshot.

The point-of-failure recovery restores the old image of the data volumes only and then replays the logs to accomplish a roll-forward recovery, provided that the log volume is intact.

Note: After completing a point of failure (POF) recovery of a single database, Veritas recommends using the VSS Snapback wizard to reattach and resynchronize the other databases and to use the VSS Snapshot wizard to create a new snapshot set.

For Exchange 2007, there is also an option to restore to a Recovery Storage Group (RSG) rather than to the production store.

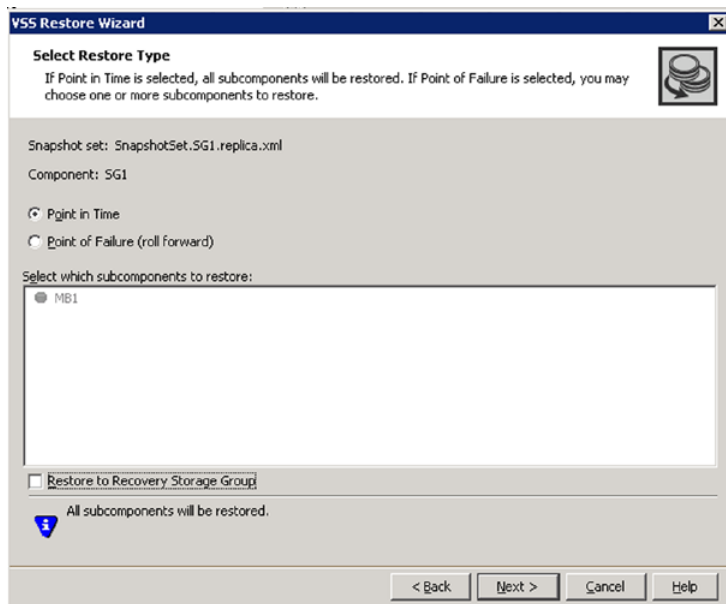
Note: COW snapshots are not supported when restoring to a recovery storage group (RSG).

Restoring the snapshot set (Exchange 2007)

The following is the procedure to restore the point in time (PIT) of the snapshot set.

To restore to the point in time of the snapshot set

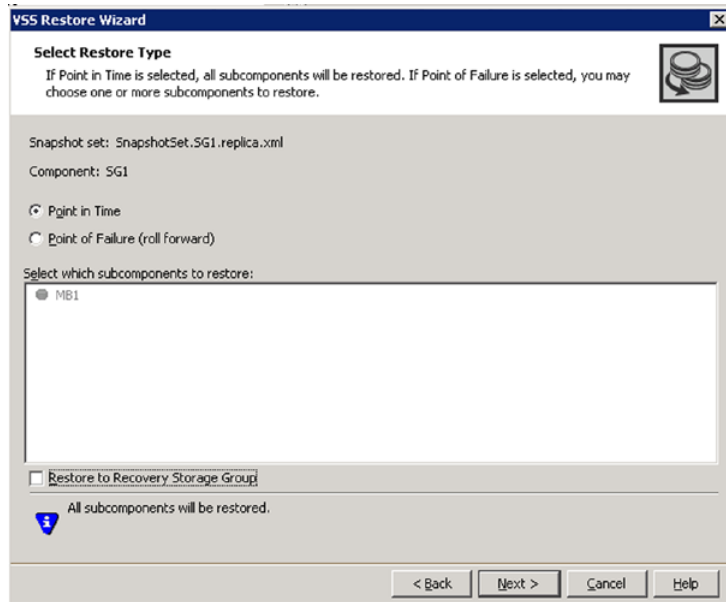
- 1 Close the database application VEA GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the volumes.
- 2 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 3 Expand the system icon, Storage Agent node, and the **Applications** node.
- 4 Right-click **Exchange** and click **VSS Restore**.
- 5 Review the **Welcome** page and click **Next**.



- 6 Select the snapshot set you want to restore and click **Next**.

The XML metadata file contains all required information that is needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 In the **Select Restore Type** panel, select Point in Time. Point in Time restores to the point in time that the snapshot set was created or refreshed based on the information that is contained in the metadata file.



Click **Next**

- 8 Verify that the restore specifications are correct and click **Finish**.
 - 9 Mount all the databases (stores) in the Exchange storage group.
 - 10 To refresh the snapshot set use the VSS Snapshot Wizard or `vxsnap create` command to create a new snapshot set of all the volumes in the storage group.
- To perform a roll-forward recovery of one or more databases to the point of failure**
- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the volumes.
 - 2 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
 - 3 Expand the system icon and the **Applications** node.
 - 4 Right-click **Exchange** and click **VSS Restore**.
 - 5 Review the Welcome page and click **Next**.

- 6 Select the name of the metadata file for the snapshot set you want to restore and click **Next**.

The XML metadata file contains all required information about the Exchange storage group, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 In the **Select Restore Type** panel, select **Point of Failure** and select one or more subcomponents (databases). Do not select the transaction logs volume. Click **Next**.

- 8 Verify that the restore specifications are correct and click **Finish**.

- 9 Mount all databases (stores) in the Exchange storage group; the transaction logs rolls forward to the point of failure. This action can be observed in the Event Viewer Application log.

- 10 To refresh the snapshot set, snapback any remaining database volumes and the transaction logs volume using either the VSS Snapback wizard or the vxsnap reattach command. Then use the VSS Exchange Snapshot wizard or vxsnap create command to create a new snapshot set of the storage group.

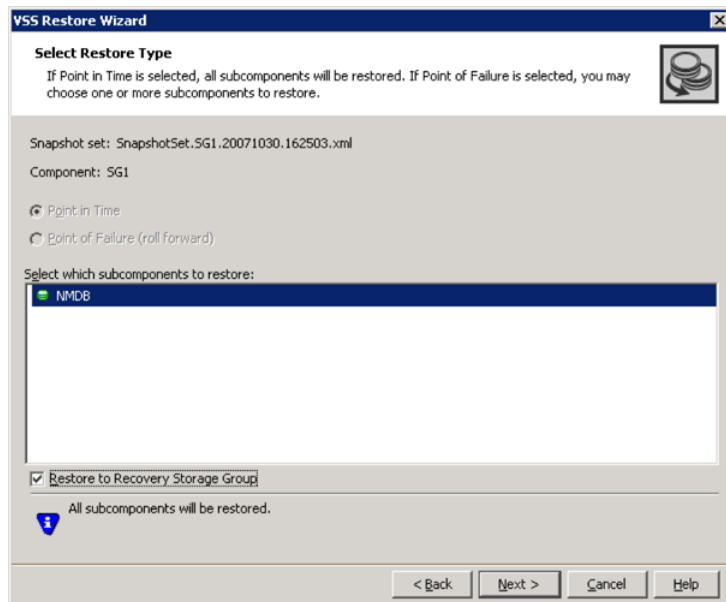
To restore to a recovery storage group (Exchange 2007 only)

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the volumes.
- 2 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 3 Expand the system icon and the **Applications** node.
- 4 Right-click Exchange and click **VSS Restore**.
- 5 Review the Welcome page and click **Next**.

- 6 Select the snapshot set you want to restore and click **Next**.

The XML metadata file contains all required information that is needed to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by File Name or Creation Time.

- 7 In the **Select Restore Type** panel, select Restore to Recovery Storage Group and select one or more subcomponents.



Click **Next**.

Restoring the snapshot set (Exchange 2010)

The VSS Restore Wizard helps you to restore a VSS shadow copy of the mailbox database components using the snapshot set files. Mailbox database can be restored with the Restore Wizard. Before you proceed with the restore operation, ensure that you have prepared and created a snapshot set with VSS Snapshot Wizard or the vxsnap utility.

After completing the restore operation on a snapshot set file, you can refresh it by using either the VSS Snapback wizard or vxsnap reattach command to snapback any remaining database volumes. Then, you can use the VSS Exchange Snapshot Wizard or vxsnap create command to create a new snapshot set of the mailbox database.

SFW automatically dismounts and mounts the stores before and after the restore operation respectively. If automatic dismount fails, then the restore operation fails. In that case, you may need to manually dismount the stores. For a VCS cluster setup, it offlines and onlines the resource database before and after the restore operation. If automatic offlining fails, then manual intervention is required.

Specifying the `-a` option in `vxsnap restore` command automatically dismounts and mounts the stores and in case of a VCS cluster setup, it offlines and onlines the VCS resource database.

For an Exchange 2010 database that is configured under VCS cluster, offline the corresponding Exchange 2010 database VCS resource before the restore operation if automatic offlining of resource database fails. Once the restore operation completes successfully, online the VCS resource.

To restore Point in Time (PIT) snapshot set, perform the following steps:

To restore Point in Time (PIT) snapshot set

- 1 Launch Exchange Management Console from **Start > All Programs > Microsoft Exchange Server 2010 > Exchange Management Console** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Expand the **Organization Configuration** node, and then select **Mailbox**.
- 2 If replication is enabled for the selected database, you should first suspend continuous replication on this database. In the result pane, on the **Database Management** tab, select the database that you want to suspend.
- 3 In the work pane, on the **Database Copies** tab, right-click the database for which you want to suspend continuous replication. Select **Suspend Database Copy** from the context menu. Click **Yes** to suspend continuous replication.
- 4 From the work pane of the snap-in, right-click on mailbox database and select **Properties > Maintenance**. Select the “**This database can be overwritten by a restore**” check box.
- 5 Now proceed with the VSS restore operations on the selected mailbox database.
Close the database application GUI, all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the snapshot volumes.
- 6 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 7 Expand the system icon and the **Applications** node.
- 8 Right-click **Exchange** and select **VSS Restore** from the context menu.
- 9 Review the Welcome page. Click **Next**.

- 10 Select the snapshot set you want to restore and click **Next**.

The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction log volumes.

Click the appropriate header to sort the list of available files by File Name or Creation Time.

- 11 In the **Select Restore Type** panel, select **Point in Time**.

Point in Time restores to the point-in-time that particular snapshot set was created or refreshed based on the information in the metadata file.

PIT restore Exchange 2010

Click **Next**.

- 12 Verify that the restore specifications are correct on the summary page and Click **Finish** to exit the wizard.
- 13 If you had suspended the replication for your clustered mailbox database as shown in step 2, update it now from the work pane of Exchange Management Console. In the work pane, on the **Database Copies** tab, right-click the mailbox database and select **Update Database Copy**.

The point-of-failure recovery restores the old image of the data volumes only and then replays the logs to accomplish a roll-forward recovery, provided that the log volume is intact.

To do a roll forward recovery to the Point of Failure (POF) of mailbox databases, perform the following:

Note: Do not enable the circular logging option. If circular logging is enabled, you cannot roll forward a database to achieve a point-of-failure recovery. Microsoft does not recommend circular logging in production environments. Refer to the Microsoft Exchange Server documentation for further details.

Note: For an Exchange 2010 database that is configured under VCS cluster, offline the corresponding Exchange 2010 database VCS resource before the restore operation. Once the restore operation completes successfully, online the VCS resource.

In other configurations, it dismounts the databases before the restore operation and then mounts the databases after the restore operation.

To perform a roll-forward recovery to the point-of-failure (POF)

- 1** Launch Exchange Management Console from **Start > All Programs > Microsoft Exchange Server 2010 > Exchange Management Console** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Expand the **Organization Configuration** node, and then select **Mailbox**.
- 2** If replication is enabled for the selected database, you should first suspend continuous replication on this database. In the result pane, on the **Database Management** tab, select the database that you want to suspend.
- 3** In the work pane, on the **Database Copies** tab, right-click the database for which you want to suspend continuous replication. Select **Suspend Database Copy** from the context menu. Click **Yes** to suspend continuous replication.
- 4** From the work pane of the snap-in, right-click on mailbox database and select **Properties > Maintenance**. Enable the “**This database can be overwritten by a restore**” check box.
- 5** Now proceed with the VSS restore operations on the selected mailbox database. Close the database application GUI, all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the snapshot volumes.
- 6** From the VEA console, navigate to the system where the production volumes and snapshot set files are located.
- 7** Expand the system icon and the **Applications** node.
- 8** Right-click **Exchange** and select **VSS Restore** from the context menu.
- 9** Review the information on the Welcome page and click **Next**.
- 10** Select the name of the XML metadata file for the snapshot set you want to restore and click **Next**. The XML metadata file contains all required information about the Exchange storage group, including the names of the database and transaction log volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 11** In the **Select Restore Type** Panel, select **Point of Failure (roll forward)**. Roll forward restores database to the most recent state before a failure occurred. Click **Next**.
- 12** Verify the restore specifications on the summary page and click **Finish** to exit the wizard.
- 13** If you had suspended the replication for your clustered mailbox database in step **2**, update it now from the work pane of Exchange Management Console. In the work pane, on the **Database Copies** tab, right-click the mailbox database and select **Update Database Copy**.

Restoring the database (POF)

The following is the procedure to restore the database to the point of failure (POF).

To perform a roll-forward recovery of one or more databases to the point of failure

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the volumes.
- 2 From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 3 Expand the system icon, Storage Agent node, and the **Applications** node.
- 4 Right-click **Exchange** and click **VSS Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the name of the metadata file for the snapshot set you want to restore and click **Next**.

The XML metadata file contains all required information about the Exchange mailbox database, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by File Name or Creation Time.

- 7 In the **Select Restore Type** panel, select **Point of Failure** and select one or more subcomponents (databases). Do not select the transaction logs volume. Click **Next**.
- 8 Verify that the restore specifications are correct and click **Finish**.
- 9 Mount all mailbox databases in the Exchange; the transaction logs roll forward to the point of failure. This action can be observed in the Event Viewer Application log.
- 10 To refresh the snapshot set, snapback any remaining database volumes and the transaction logs volume using either the VSS Snapback wizard or the `vxsnap reattach` command. Then use the VSS Snapshot wizard or `vxsnap create` command to create a new snapshot set of the storage group.

Restoring to a recovery storage group

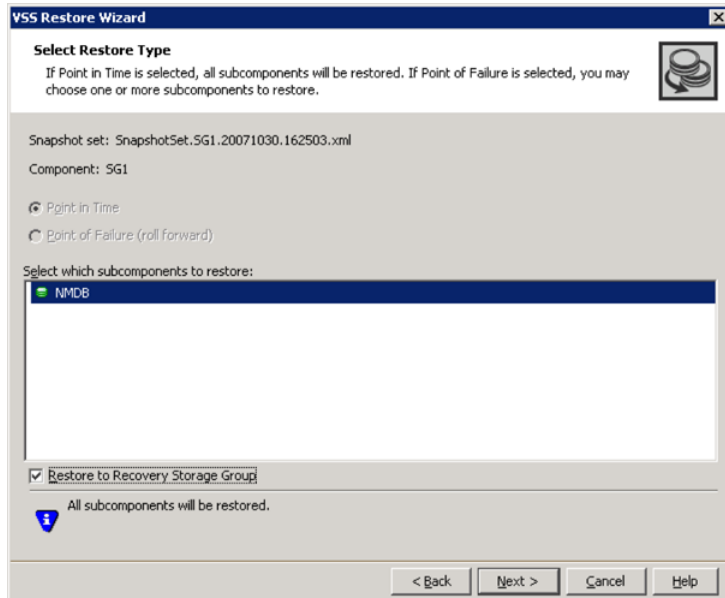
The following is the procedure to restore Microsoft Exchange to a recovery storage group.

To restore to a recovery storage group (Exchange 2007 only)

- 1** Close the database application GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the volumes.
- 2** From the VEA console, navigate to the system where the production volumes and snapshot set are located.
- 3** Expand the system icon, Storage Agent node, and the **Applications** node.
- 4** Right-click **Exchange** and click **VSS Restore**.
- 5** Review the Welcome page and click **Next**.
- 6** Select the snapshot set you want to restore and click **Next**.

The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 In the **Select Restore Type** panel, select **Restore to Recovery Storage Group** and select one or more subcomponents.



Click **Next**

Note: The **Restore to Recovery Storage Group** option is not supported for restoring from COW snapshots.

- 8 Verify that the restore specifications are correct and click **Finish**.
- 9 Mount all the databases (stores) in the Exchange storage group.

Refreshing the snapshot set

The VSS Refresh command available in the context menu of the VSS Writer object refreshes the GUI display of the VSS Writer and components. It does not refresh the snapshot set.

To refresh the snapshot set, either use the **VSS Snapback Wizard** for a one-time refresh or set up a schedule using the **VSS Snapshot Scheduler Wizard**.

Supporting replication in Exchange 2007

SFW provides support for snapshots when the Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR) features of Exchange 2007 are enabled.

The following topics describe various notes about support of replication in Exchange 2007.

VEA GUI

If replication for Exchange 2007 is enabled, then the display of the Microsoft Exchange Replication Service instance of the Microsoft Exchange Writer is enabled and displayed in the VEA GUI. It appears as the Microsoft Exchange Writer Replica and is displayed in the tree view of the VEA subordinate to the VSS Writers node.

Right-clicking Exchange Replica node displays a context menu that shows VSS Refresh, VSS Snapshot, Schedule VSS Snapshot, VSS Snapback, and Properties. Restoring the replica with the VSS Restore operation is not supported.

Note: The Prepare command is required before using VSS Snapshot.

Microsoft Failover Clustering environment

In a Microsoft Failover Clustering environment, you have to manually set the dependency of the Microsoft Exchange database instance to the Volume Manager Disk group resource so that it fails over in the correct sequence.

vxsnap CLI command

SFW provides a `vxsnap` CLI command option to take a snapshot of a Microsoft Exchange Replica (Microsoft Exchange Replication Service instance of the Microsoft Exchange writer) or of a Microsoft Exchange Writer (Microsoft Exchange Service instance of the Microsoft Exchange writer).

In the command, you can specify the replica store writer to take the snapshot of the replica or the active store writer to take the snapshot of the active store. If the replica store writer or the active store writer is not specified, then "Microsoft Exchange Writer" is used as a default.

For example:

```
vxsnap -x snapdata.xml create writer="Microsoft Exchange Replica"  
component=SG1 backupType=COPY -E -O
```

specifies that the VSS Writer, Microsoft Exchange Replica, is used to take a snapshot of the replica.

Note: The Prepare operation must be completed on the volumes that contain the replica before taking a snapshot of a replica. This can be done using the VEA GUI or the `vxsnap prepare` CLI command. When using the CLI, the `vxsnap prepare` command must specify the Microsoft Exchange Replica.

For example:

```
vxsnap prepare component=SG1/writer="MicrosoftExchange Replica"  
source=L:/harddisk=harddisk2
```

See [“vxsnap CLI command”](#) on page 320.

Using a Snapshot of a Replica for Database Recovery

A snapshot of a replica can be used to restore an Exchange database to the point of failure (POF) or be used to restore a complete Exchange mailbox to a point in time (PIT).

Taking a snapshot of a replica can be done with the GUI by using the VSS snapshot wizard. Taking a snapshot of a replica automatically takes a snapshot of all the volumes that the replica uses.

Originally, a manual dismount of the database was required before the restore operation. Now SFW performs the dismount of the database automatically as a part of the restore operation.

Note: In a VCS environment, the database is automatically dismounted and the database is set for overwrite by restore as a part of the restore operation. However in a Microsoft Failover Clustering environment, a manual dismount of the database and manually setting the database for overwrite by restore are both required.

Note: When SFW fails to automatically dismount the database during a restore operation, the restore operation fails. The restore operation can be performed again after manually dismounting the databases and manually setting the databases for overwrite by restore. If LCR is enabled in a non-Microsoft Failover Clustering environment or if CCR is enabled in a Microsoft Failover Clustering environment, suspend LCR/CCR. Then repeat the restore operation. If you suspended LCR/CCR, re-enable it after the restore operation.

To perform the restore on the active writer, use the Exchange Management shell to execute specific cmdlets.

These cmdlets are the following:

- Dismount Database cmdlet

```
dismount-Database -Identity <DatabaseIdParameter>  
[-DomainController <Fqdn>]
```

For more information about the Exchange Management shell and cmdlets, refer to the Microsoft Exchange 2007 product information at

<http://www.microsoft.com>.

Note: Although SFW lets you restore the database from a snapshot, restoring just the database log files is not supported.

An example of a PIT recovery procedure from a snapshot of a replica of an Exchange storage group, SG1, that contains two databases, DB1 and DB2, on an Exchange server, TestExch, would be as follows.

To perform a PIT recovery procedure (example):

1 Run Dismount Database cmdlet on DB1 and DB2 databases.

```
dismount-Database -Identity TestExch\SG1\DB1  
dismount-Database -Identity TestExch\SG1\DB2
```

2 Run RestoreStorageGroupCopy cmdlet on SG1 storage group.

```
Restore-StorageGroupCopy -Identity TestExch\SG1 -ReplaceLocations
```

3 Run Mount Database cmdlet on DB1 and DB2 databases.

```
mount-Database -Identity TestExch\SG1\DB1  
mount-Database -Identity TestExch\SG1\DB2
```

4 Perform refresh.

```
vxsnap refresh
```

5 Perform VSS restore operation using snapshot of replica.

```
vxsnap -x snapdata.xml restore RestoreType=PIT  
writer="Microsoft Exchange Writer"
```

For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create  
writer="Microsoft Exchange Replica"  
component=SG1 backupType=COPY-E -O
```

Additional considerations for a Microsoft Failover Clustering environment

For a Microsoft Failover Clustering environment, there are additional considerations when restoring the database.

These include the following:

- After performing the refresh operation, the user must manually dismount the databases before performing the restore operation.
- If CCR is enabled, then the user must disable the circular logging feature and suspend CCR before restoring the database.

Supporting replication in Exchange 2010

If replication for Exchange 2010 is enabled, then display of Microsoft Exchange Writer is enabled and displayed in the VEA GUI as Exchange Replica. It is displayed in the tree view of the VEA subordinate to the Applications node.

Right-clicking Exchange Replica node displays a context menu that shows VSS Refresh, VSS Snapshot, Schedule VSS Snapshot, VSS Snapback, and Properties. Restoring the replica with the VSS Restore operation is not supported.

The Prepare command is required before using VSS Snapshot.

Using a snapshot of a replica for mailbox database recovery in Database Active Group (DAG)

A snapshot of a replica can be used to restore an Exchange mailbox database to the point of failure (POF) or to a point in time (PIT). Taking a snapshot of a replica can be done through the VEA GUI by using either the VSS snapshot wizard or with

the vxsnap CLI command. Taking a snapshot of a replica automatically takes a snapshot of all the volumes that the replica uses.

When SFW fails to automatically dismount the database during a restore operation, the restore operation fails. The restore operation can be performed again after manually dismounting the databases and manually setting the databases for overwrite by restore.

Although SFW lets you restore the database from a snapshot, restoring just the database log files is not supported.

For an Exchange 2010 database that is configured under VCS cluster, offline the corresponding VCS database resource before the restore operation if automatic offlining of resource database fails. SFW automatically offlines the resource databases. Once the restore operation completes successfully, online the VCS resource database.

In other configurations, SFW dismounts the databases before the restore operation and then mounts the databases after the restore operation automatically.

Using the `-a` option automatically dismounts and mounts the databases. For a VCS cluster setup, it offlines and onlines the resource databases.

Refer to Microsoft Exchange 2010 Shell Help for details on Exchange cmdlets.

An example of a PIT recovery procedure from a snapshot of a replica copy of an Exchange mailbox database **DB1** on an active mailbox database **TestServer1** and Replica copy on **TestServer2** would be as follows:

- Take Snapshot of the replica copy of mailbox database on **TestServer2**.
- Use the `Move-ActiveMailboxDatabase` cmdlet to perform the active mailbox database **DB1** switchover from **Test Server 1** to **Test Server 2**.

```
Move-ActiveMailboxDatabase MB1 -ActivateOnServer TestServer2
```

Now **TestServer2** has the Active mailbox database and **TestServer1** has Replica copy.

- Use the `Suspend-MailboxDatabaseCopy` cmdlet to block replication and replay activities (log copying and replay)

```
Suspend-MailboxDatabaseCopy -Identity DB1\TestServer1 [-Confirm:$False]
```

Use `-Confirm` to prevent user intervention

- Use the `Set-MailboxDatabaseCopy` cmdlet to configure the properties of a database copy and set check box for “**This database can be overwritten by Restore.**”

```
set-mailboxdatabase -identity MB1\TestServer2 -allowfilerestore:$true
```


- Use the `Dismount-Database` cmdlet to dismount the database

```
dismount-Database -Identity MB1\TestServer2 [-Confirm:$False]
```

Use `-Confirm` to prevent user intervention

For a VCS cluster setup, offline the Exchange 2010 VCS resource database if automatic offlining fails.

- Now perform the VSS restore operation using snapshot of replica.

```
vxsnap -x snapdata.xml restore RestoreType=  
PIT writer="Microsoft Exchange Writer"
```

For this example, assume that the snapshot of the replica was performed with

```
vxsnap -x snapdata.xml create writer=  
"Microsoft Exchange writer Replica" component=DB1 backupType=COPY -E -O
```

- Use the `Mount-Database` cmdlet to mount database DB1

```
mount-Database -Identity DB1\TestServer2
```

For a VCS cluster setup, online the Exchange 2010 VCS resource database if automatic onlining of resource database fails.

- Use the `Update-MailboxDatabaseCopy`

```
Update-MailboxDatabaseCopy  
-Identity MB3\TestServer2 -SourceServer TestServer1 -DeleteExistingFiles [
```









Use `-Confirm` to prevent user intervention

Using the VSS snapshot wizards with Enterprise Vault

SFW provides support for taking snapshots of Enterprise Vault components. FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all volumes that are associated with a Enterprise Vault component without taking the component's databases offline. The Enterprise Vault Snapshot wizard uses VSS to take snapshots of components. The VSS process allows the component's databases to be quiesced before the snapshot operation and reactivates them immediately after. The quiescing at the component level allows for supported and guaranteed persistent snapshots of the data. A snapshot of a component can be reattached and resynchronized to match the current state of the component with the Enterprise Vault Snapback wizard. SFW also provides recovery support for Enterprise Vault components. Using the Enterprise

Vault Restore Wizard, the snapshots that are taken with the Enterprise Vault Snapshot wizard can be used for a recovery of the component.

The following describes the Enterprise Vault components that SFW supports.

	Directory database	The Enterprise Vault Directory database holds configuration information for each site in the Directory (EnterpriseVaultDirectory).
	Monitoring database	Enterprise Vault has a Monitoring agent on each Enterprise Vault server. The Monitoring agents collect monitoring data at scheduled intervals, usually every few minutes, and store it in the Enterprise Vault Monitoring database. If multiple Enterprise Vault sites share a Directory database, then they must also share a Monitoring database (EnterpriseVaultMonitoring).
	Reporting database	The Reporting database stores reporting information that is generated from the collected Monitoring data. There is one Reporting database named EnterpriseVaultFSAReporting corresponding to the Directory Database.
	Auditing database	The Auditing database stores audit information. There is one Auditing database named EnterpriseVaultAudit corresponding to the Directory database
	Site	A site comprises one or more Enterprise Vault servers running one or more Enterprise Vault services and tasks to archive items from specified targets. A site also contains a collection of vault stores, archiving policies that define how and when items are to be archived, and Retention Categories that define how long items are to be stored before being deleted.
	Index	The Indexing Service indexes items as they are archived. There is one Index for each archive
	Vault Store Group	A Vault Store Group (VSG) is a collection of Vault Stores that are configured to share data.
	Fingerprint database	A Vault Store Group contains one or more Fingerprint Catalog databases (FDB) to hold a shared item's metadata.

**Vault Store**

Enterprise Vault organizes archives in entities called Vault Stores. Vault Stores contain one or more Enterprise Vault Partitions. A Partition can reside on any of the supported storage media.

**Vault Store database**

The Vault Store SQL databases contain information about every archive operation processed for all archives within a Vault Store. Both the Vault Store's database and the Vault Store's files must be complete to be able to retrieve messages from an archive.

**Open Partition**

In each Vault Store, there can be only one Open Partition. This is the Partition in which Enterprise Vault archives data. An Open Partition is the current active Partition in which Enterprise Vault data is being written.

**Ready Partition**

Each Vault Store can have a number of Ready Partitions. A Ready Partition is the next available partition to be used.

**Closed Partition**

Each Vault Store can have any number of Closed Vault Store Partitions. Enterprise Vault does not archive content into Partitions that are closed.

Using the Enterprise Vault Snapshot wizard

Creating a snapshot is a two-step process. The first step, prepare the volume for snapshot, creates snapshot mirrors attached to all the original volumes of the specified component. Depending on the size and number of volumes, the process of synchronizing the new snapshot mirrors with the original production volumes can take a long time. The second step uses the Enterprise Vault Snapshot wizard to create the snapshot set by detaching the snapshot mirrors from the original volumes and creating separate on-host snapshot volumes as well as an XML file to store the component and snapshot volume metadata. The Enterprise Vault Snapshot wizard integrates with VSS to quiesce the component and then simultaneously snapshot the volumes in the component. This snapshot is done while the component is online and without disrupting processing. Once a snapshot set has been created, it can be reattached and resynchronized with the Enterprise Vault Snapback wizard.

Note: After performing operations for snapshots, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

Preparing for Enterprise Vault snapshots

After completing Enterprise Vault installation tasks, and before taking Enterprise Vault snapshots, you must invoke the Enterprise Vault Management Shell to register PowerShell snap-ins, set the execution policy and then configure the Veritas Scheduler Service. The Enterprise Vault Management Shell can be invoked either before or after performing the Prepare operation for snapshot mirrors.

Setting the execution policy with the Enterprise Vault Management Shell

You must invoke the Enterprise Vault Management Shell to set the execution policy to unrestricted.

You need to invoke the Enterprise Vault Management Shell only once on a host where you perform the snapshot command. As a best practice, it is recommended that you invoke the shell once on all the hosts in the Enterprise Vault configuration.

To set the execution policy with the Enterprise Vault Management Shell

- 1 Launch Enterprise Vault Management Shell from **Start > All Programs > Enterprise Vault > Enterprise Vault Management Shell** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Run the following command to set the execution policy:

```
set-executionpolicy unrestricted
```

Caution: Setting the execution policy to unrestricted may be a potential security risk. This setting allows any (untrusted) script to be invoked.

Configuring the Veritas Scheduler Service

Configure the Veritas Scheduler Service with administrative privileges on all nodes where you have configured the Enterprise Vault server.

Perform the following to configure the service.

To configure the Veritas Scheduler Service

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

Under **Services**, select **Veritas Scheduler Service**.
- 2 From the Veritas Scheduler Service Properties dialog box, perform the following settings:
 - Select **Log On**.

- Disable the **Local System account** option.
 - Enable the **This account** option.
 - Ensure that the user name is set up with administrative privileges.
- 3 Click **OK** to complete the configuration.

Preparing an Enterprise Vault snapshot mirror

Preparing Enterprise Vault (EV) for a snapshot involves preparing all the dependent components of Enterprise Vault. The snapshot prepare operation can be performed at the Enterprise Vault farm, site, VaultStore Group, VaultStore, or the component level. When you right-click an Enterprise Vault object and select the Enterprise Vault Prepare operation from the menu option, related Enterprise Vault sites, components, etc. within (subordinate to) the Enterprise Vault object are prepared for snapshot mirrors.

Note: You can have only one snapshot volume on a disk for particular production volume.

Note: If a volume is used for more than one component, it is prepared for snapshot only once when all the components are specified in a vxsnap prepare CLI command or in the Enterprise Vault Prepare wizard. If several components share the same volume and these components are each prepared separately for snapshot mirrors, then the volume is prepared many times (once for every component).

The snapshot mirrors remain attached to the original volumes and continue to be updated until you use the Enterprise Vault Snapshot Wizard, the vxsnap create command, or the Enterprise Vault Snapshot Scheduler Wizard to create the snapshot set.

Please note the following:

- When an Enterprise Vault site is Prepared, the Directory database is also Prepared implicitly by the system along with the Enterprise Vault site components.
- When a VaultStore database of a VaultStore is prepared, the Open Partition of that VaultStore is also prepared implicitly by the system.
- When an Open Partition of a VaultStore is prepared, the VaultStore database of that VaultStore is also prepared implicitly by the system.
- Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, %VMPATH%\logs\EVstatus.log. The log

contains information about the success or failure of the operation for the components. In the event that the prepare of a volume for a component fails, the operation continues to prepare the remaining volumes of the component and any other requested components. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

You can also use the `vxsnap prepare` CLI command to prepare volumes as snapshot mirrors.

To prepare an Enterprise Vault snapshot mirror

- 1 On the VEA console, navigate to the Storage Agent in the tree-view and expand the Enterprise Vault nodes, under **Applications**.
- 2 Right-click the desired Enterprise Vault node and click **Enterprise Vault Prepare** on the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.
- 4 In the **Select Component** dialog, select the Enterprise Vault components (sites, VaultStoreGroups, VaultStores, or partitions) that you want to snapshot prepare.

Review the specifications for the prepare operation and click **Finish**.

Note: Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each node. This is applicable for SQL Server 2012 or later.

Creating the snapshot set

SFW provides support for creating snapshots for the Enterprise Vault (EV) farm, sites, VaultStore Groups, VaultStores, and components. The Enterprise Vault Snapshot Wizard integrates with Microsoft Volume Shadow Copy Service (VSS) to let you create snapshots of all volumes associated with Enterprise Vault components and volumes without taking the databases offline and disrupting database operations. It also allows the Enterprise Vault Powershell cmdlet to put Enterprise Vault into backup mode. Backup mode ensures that there are no write operations involving any Enterprise Vault components and helps avoid disruption.

The resulting snapshot set provides a complete picture of the database at a point in time the snapshot was taken. You can take a snapshot with the Enterprise Vault Snapshot Wizard or the `vxsnap create` CLI command.

Note: The Enterprise Vault Snapshot Wizard and vxsnap utility can be invoked from any node that has the Directory service running.

The snapshot operation can be performed at the Enterprise Vault farm, site, VaultStore Group, VaultStore, or component levels. When you right-click an Enterprise Vault object and click Enterprise Vault Snapshot from the context menu, the Enterprise Vault objects and components within (subordinate to) the selected Enterprise Vault object are all included in the snapshot. You can also add more Enterprise Vault objects from the objects shown in the Enterprise Vault Snapshot wizard.

Please note the following:

- When an Enterprise Vault site is snapshotted, the Directory database is also snapshotted implicitly by the system along with the Enterprise Vault site components.
- When a Vaultstore database is snapshotted, the Open Partition of that VaultStore is also snapshotted implicitly by the system.
- When an Open Partition of a VaultStore is snapshotted, the Vaultstore database of that VaultStore is also snapshotted implicitly by the system.
- Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, %VMPATH%\logs\EVStatus.log. The log contains information about the success or failure of the operation for the components. In the event that the snapshot of a volume for a component fails, the operation does not snapshot the remaining volumes of the component nor any other requested components that are remaining. If the operation succeeds for all the volumes of a component, then the status of the component is logged as SUCCESS. If the operation fails for any one of the volumes of the component, then the status of all the components of that Enterprise Vault server is logged as FAILED along with the cause of the failure. The status of all components that had been already snapshotted earlier before the failed component are logged as SUCCESS, and the status of the remaining components that had not been snapshotted are logged as SKIPPED.

Review the following before proceeding with the snapshot operation for Enterprise Vault:

- Ensure that the execution policy is set to unrestricted. You can use the Enterprise Vault Management Shell to set the execution policy. See [“Setting the execution policy with the Enterprise Vault Management Shell”](#) on page 328.
- Ensure that the Veritas Scheduler Service is configured for all Enterprise Vault servers in the Enterprise Vault configuration.

See [“Configuring the Veritas Scheduler Service”](#) on page 328.

- Ensure that the component volumes or database and partition volumes have been snapshot prepared as snapshot mirrors.
See [“Preparing an Enterprise Vault snapshot mirror”](#) on page 329.

The following is the procedure to create a snapshot set.

To create the snapshot set

- 1** Using the VEA console, navigate to the **Storage Agent** node in the tree-view and expand the **Enterprise Vault** nodes under **Applications**.
- 2** Navigate to the desired Enterprise Vault node.
Right-click the Enterprise Vault node and click **Enterprise Vault Snapshot** in the context menu.
- 3** In the wizard, review the Welcome page and click **Next**.

4 Specify the snapshot set parameters and then click **Next**

Select Component for snapshot operation	Select the component for the snapshot set
Directory	<p>The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory shown on the screen.</p> <p>There are two ways to change the XML file location.</p> <p>First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named <code>redirect.txt</code>. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, <code>\\ServerName\SharedFolderName</code>. Save the <code>redirect.txt</code> file in the default directory, for example: <code>C:\Program Files\Veritas\Veritas Volume Manager\VSSXML</code>.</p>
Snapshot set	Enter a name for the snapshot set, for example, <code>billing</code> or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time. If multiple components are selected, "EV" is used in the default snapshot set name instead of the component name.
Select snapshot type	Select the snapshot type.
	<p>You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a Enterprise Vault component. However, if you want to replay the logs as a part of restoring a component, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.</p>

5 Review the specifications of the snapshot set and click **Finish**.

Using the Enterprise Vault Snapback wizard

The Enterprise Vault Snapback Wizard reattaches and resynchronizes an existing snapshot set so that it matches the current set of its original database.

Note: Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, %VMPATH%\logs\EVStatus.log. The log contains information about the success or failure of the operation for the components. In the event that the reattach of a volume for a component fails, the operation continues to reattach the remaining volumes of the component and any other requested components. The components that successfully complete the operation are removed from the snapshot set. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

To snapback an Enterprise Vault snapshot set

- 1 Except for the VEA GUI, close all windows, applications, or third-party system management tools that may access the snapshot set.
- 2 Using the VEA console, navigate to the **Storage Agent** node in the tree-view and expand the **Enterprise Vault** nodes under **Applications**.
- 3 Right-click the Enterprise Vault node and click **Enterprise Vault Snapback** in the context menu.
- 4 In the wizard, review the Welcome page and click **Next**.
- 5 Select the snapshot set and specify the parameters and then click **Next**.

Directory	Specify the path of the directory where the snapshot set is located.
Snapshot Sets	The XML metadata file contains all the information that is required to snapback the snapshot set, including the names of the database and transaction log volumes. Click the appropriate header to sort the list of available files by File Name or Creation Time. This file is deleted after the snapback operation has completed successfully.

- 6 On the Select Components page, select the components to be reattached or snapbacked. Click **Next** to continue.
- 7 Review the specifications of the snapback and click **Finish**.

Using the VSS Enterprise Vault Snapshot Scheduler wizard

Before you run the Enterprise Vault Snapshot Scheduler Wizard to schedule a snapshot set for a database, you must prepare a snapshot mirror for each of the volumes in the database.

See [“Preparing an Enterprise Vault snapshot mirror”](#) on page 329.

You can then use the Enterprise Vault Snapshot Scheduler Wizard to schedule the initial snapshot set and to set up the schedule for keeping it refreshed.

The scheduling of snapshots can be done at the Enterprise Vault (EV) farm, sites, VaultStore Groups, Vault Stores, Partitions, or any component level. The Schedule Enterprise Vault Snapshot Wizard can be launched from any component of Enterprise Vault configuration.

When the scheduled snapshots occur, the snapshot mirrors are detached from the original volumes, creating separate on-host snapshot volumes as well as an XML file to store the snapshot volume metadata. The scheduled process integrates with VSS to quiesce the database and then simultaneously snapshot the volumes for Enterprise Vault components. This snapshot is done while the Enterprise Vault components are online and without disrupting the operations. It also allows the Enterprise Vault PowerShell cmdlet to put Enterprise Vault into backup mode. Backup mode ensures that no writes happen on Enterprise Vault components.

To schedule a snapshot for a selected component

- 1 Using the VEA console, navigate to the **Storage Agent** node in the tree-view and expand the **Enterprise Vault** nodes under **Applications**.
- 2 Right-click the appropriate Enterprise Vault node and click **Schedule Enterprise Vault Snapshot** in the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.
- 4 On the Select Components page, select the Enterprise Vault components you want to snapshot. Select the snapshot type and give it a name.

Specify the following and click **Next** to continue.

Select component for snapshot operation

Select the EV components you want to snapshot. You can select EV Farm, sites, Vault StoreGroups, Vault Stores, partitions or any component.

Additionally, you can also use the keyboard shortcuts to select EV components.

Ctrl+ A selects the entire EV farm along with sites and components. Use Ctrl+Click to select EV components of your choice. Use Ctrl+Shift+Arrow Up, Down keys to select components located at the upper and lower hierarchies in the EV farm.

Directory

The XML file is stored by default in the directory shown on the screen.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore fails.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`.

Save the `redirect.txt` file in the default directory, for example: `C:\Program`

```
Files\Veritas\Veritas Volume  
Manager\VSSXML.
```

Note: You must not use the volume name or volume path in the `redirect.txt` file that is involved in the snapshot. If the volume name or path for the snapshot is used, then a restore fails.

Snapshot set

Enter a name for the snapshot set, for example, `billing` or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time. Example: `SnapshotSet.EV.20091001.124837`.

Note: If multiple components are selected, "EV" is used in the default snapshot set name instead of the component name.

Select snapshot type

Select the snapshot type.

You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a Enterprise Vault component.

Full Backup

- A full snapshot affects the sequence of backups and is used when replay logs are needed while restoring.

If you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup.

- Copy
A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

5 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

Name of this schedule	<p>Enter a unique name for the snapshot set schedule.</p> <p>This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the VSS writer name and a numbered suffix that increments with each schedule.</p>
Description of this schedule	<p>Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.</p>
Start Time	<p>The time of the day to begin taking snapshots.</p>
End Time	<p>The time of day to end taking snapshots.</p> <p>If a snapshot is in progress it is completed but a new one is not started after the end time.</p>
Schedule takes effect on	<p>The date on which the specified schedule takes effect. The default is the current date.</p>
Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

6 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

7 Click **Next**.

8 Review the snapshot set and schedule details and click **Finish**.

Using the Enterprise Vault Restore Wizard

SFW provides recovery support for Enterprise Vault components. Using the Enterprise Vault Restore Wizard, the snapshots created with the Enterprise Vault Snapshot wizard can be used to restore the components without logs.

Restoring the component without logs (recovery and no log replay) restores the database and the transaction log volumes. No other logs are applied. This leaves the component in an online state.

Note: Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, %VMPATH%\logs\EVStatus.log. The log contains information about the success or failure of the operation for the components. In the event that the restore of a volume for a component fails, the operation continues to restore the remaining volumes of the component and any other requested components. The components that successfully complete the operation are removed from the snapshot set. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

Recovering using snapshots without log replay

The following procedure uses the Recovery option. It restores the component database from the snapshot set volumes to the time of the snapshot set. The component database and transaction log volumes are restored but no additional transaction logs are applied.

Note: The Recovery option is applicable only for Enterprise Vault database components and not for Enterprise Vault Index and Partitions.

To recover EV components without log replay using the VEA

- 1 Except for the VEA GUI, close all windows, applications, or third-party system management tools that may access the Enterprise Vault component volumes. It is also recommended to bring the database offline.
- 2 Using the VEA console, navigate to the **Storage Agent** node in the tree-view and expand the **Enterprise Vault** nodes under **Applications**
- 3 Right-click the Enterprise Vault node and click **Enterprise Vault Restore** in the context menu.
- 4 In the wizard, review the Welcome page and click **Next**.
- 5 Select the snapshot set XML metadata file to be used for this operation and click **Next**. The XML metadata file contains all the information that is required to restore the snapshot set. Click the appropriate header to sort the list of available files by File Name or Creation Time.
- 6 On the Select Enterprise Vault components for Restore operation page, select the component or components that you want to restore using the Recovery option. The original volumes of the selected components are restored from the snapshot volumes.

- 7 On the **Select Restore Type** panel, select **Recovery** and specify options.

Recovery with missing original volumes	Enable this check box if one or more original volumes are missing.
Forcefully close any open handles	Enable this check box if any Enterprise Vault component volumes have open handles. This option forcefully closes any open handles of the volumes to be restored.

Click **Next** to continue.

- 8 Verify the snapshot restore specifications on the summary page and click **Finish**.

The database is restored to the time the snapshot set was created or last refreshed. Bring the database back online, if needed.

Restoring snapshots and manually applying logs

The following procedure uses the No Recovery option to restore the component database from the component database and log snapshot volumes. Selecting this option leaves the component database in a loading state. You can then manually apply backed up transaction logs to recover the component database to the desired point in time.

Note: The No Recovery option is applicable only for Enterprise Vault database components and not for Enterprise Vault Index and Partitions.

Warning: Before you begin, use your preferred method to backup the transaction logs within SQL Server. You must use the “overwrite existing media” option to create uniquely-named backup files.

To restore EV components using the No Recovery option

- 1 Ensure that you have backed up the transaction logs within the SQL Server using the “overwrite existing media” option to create uniquely-named backup files.
- 2 Except for the VEA GUI, close all windows, applications, or third-party system management tools that may access the Enterprise Vault component volumes. It is also recommended to bring the database offline.
- 3 Using the VEA console, navigate to the **Storage Agent** node in the tree-view and expand the **Enterprise Vault** nodes under **Applications**

- 4 Right-click the Enterprise Vault node and click **Enterprise Vault Restore** in the context menu.
- 5 In the wizard, review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**. The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction log volumes. Click the appropriate header to sort the list of available files by File Name or Creation Time.
- 7 On the Select Enterprise Vault components for Restore operation page, select the Enterprise Vault component or components that you want to restore. The original volumes of selected components are restored from the snapshot volumes.
- 8 On the **Select Restore Type** panel, select **No Recovery** and specify any options.

Forcefully close any open handles

Enable this check box if any Enterprise Vault component volumes have open handles. This option forcefully closes any open handles of the volumes to be restored.

Click **Next** to continue.

- 9 Verify the snapshot restore specifications on the summary page and click **Finish**.

The component database and log snapshot volumes are restored and the component database is left in a loading state.
- 10 Use your preferred method to manually restore the backup transaction logs to the desired point in time and then bring the component database back online.

Using the VSS snapshot wizards with Microsoft SQL

SFW provides support for taking snapshots of Microsoft SQL databases. FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all volumes that are associated with an SQL database without taking the database offline. The VSS Snapshot wizard uses VSS to take snapshots of the database. The VSS process allows the database to be quiesced before the snapshot operation and reactivates it immediately after. The quiescing of the database and guarantees persistent snapshots of the data. A snapshot of a database can be reattached and resynchronized to match the current state of the database with the VSS Snapback wizard.

Note: Snapshots can only be taken of read-write volumes when VSS is used. The resulting VSS snapshot is read only. The file system label of VSS snapshot volumes cannot be changed.

SFW also provides a VSS Snapshot Scheduler wizard that can be used as an alternative to the VSS Snapshot wizard and the VSS Snapback wizard. It enables you to set up a schedule for taking the initial snapshots and for automating the snapback refresh process. At the scheduled time for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again. The schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

SFW also provides recovery support for a SQL database. Using the VSS Restore Wizard, the snapshots taken with the VSS Snapshot wizard can be used for a recovery of the database with or without logs.

Refer to the *Storage Foundation and High Availability Solutions Solutions Guide for Microsoft SQL* for additional information about how to use FlashSnap with Microsoft SQL to perform and to implement recovery procedures.

See [“Using the VSS Snapshot wizard”](#) on page 343.

See [“Using the VSS Snapback wizard”](#) on page 347.

See [“Using the VSS Snapshot Scheduler wizard”](#) on page 347.

See [“Using the VSS Restore Wizard ”](#) on page 355.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support snapshot or VSS-related commands.

Using the VSS Snapshot wizard

Creating a snapshot is a two-step process. The first step, prepare the volume for snapshot, creates snapshot mirrors attached to all the original volumes in the specified database. Depending on the size and number of volumes, the process of synchronizing the new snapshot mirrors with the original production volumes can take a long time. The second step uses the VSS Snapshot wizard to create the snapshot set by detaching the snapshot mirrors from the original volumes and creating separate on-host snapshot volumes as well as an XML file to store the database and snapshot volume metadata. The VSS Snapshot wizard integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. This snapshot is done while the database is online and without disrupting processing. Once a snapshot set has been created, it can be reattached and resynchronized with the VSS Snapback wizard.

Note: After performing operations for snapshots, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

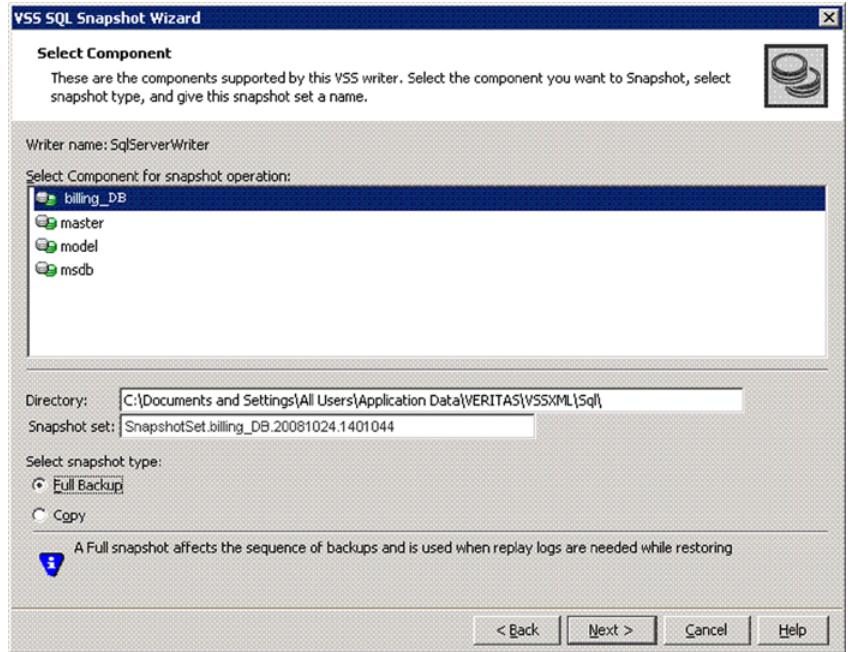
Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each node. This is applicable for SQL Server 2012 or later.

Creating the snapshot set

The following is the procedure to create a snapshot set.

To create the snapshot set

- 1 From the VEA console, navigate to the system where the production volumes and snapshots mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
If the Applications node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQL** node.
- 4 Right-click the instance and click **VSS Snapshot**.
- 5 In the wizard, review the Welcome page and click **Next**.

6 Specify the snapshot set parameters as follows and then click **Next**:

Select Component Select the database for the snapshot set.
for snapshot
operation

Directory

The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory that is shown on the screen.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`. Save the `redirect.txt` file in the default directory, for example: `C:\Program Files\Veritas\Veritas Volume Manager\VSSXML`.

Snapshot set	Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time.
Select snapshot type	Select the snapshot type. You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

- 7 In the **Change Attributes** panel, optionally change the attributes for the snapshot volumes and click Next:

Volume Name	Displays the name of the volume.
Snapshot Volume Label	Displays the read-only label for the snapshot volume.
Drive Letter	Optionally, click a drive letter and select a new choice from the drop-down menu.
Plex	Optionally, click a plex and select a new choice from the drop-down menu.

- 8 On the **Synchronized Snapshot** panel (Volume Replicator only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Volume Replicator (Volume Replicator). Otherwise, you are directly taken to the **Schedule Information** panel.

See *Volume Replicator Administrator's Guide*.

- 9 Review the specifications of the snapshot set and click **Finish**.

Using the VSS Snapback wizard

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot set so that it matches the current state of its original SQL database. The wizard is available in the context menu of the VSS Writer object.

Resynchronizing the snapshot set

The following is the procedure to resynchronize (snapback) a snapshot set.

To snapback a snapshot set

- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may access the snapshot set.
- 2 From the VEA, select the system where the production volumes and snapshot mirrors are located, as the active host.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click the node of the selected application and click **VSS Snapback**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set you want to snapback and click **Next**.

The XML metadata file contains all the information that is required to snapback the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**. This file is deleted after the snapback operation has completed successfully.

- 7 If a message appears that indicates some volumes have open handles, confirm that all open handles are closed and then click Yes to proceed.
- 8 Verify that the snapback specifications are correct and click **Finish**.

Using the VSS Snapshot Scheduler wizard

You can use the VSS Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of refreshing snapshots sets. At the scheduled time for the snapshot, the snapshot volumes are automatically reattached, resynchronized, and then split again. Once configured and applied, the schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background. In a clustered server environment, ensure that the scheduler service is configured on each node with domain administrator privileges so that any node in the cluster can run the schedule after a failover.

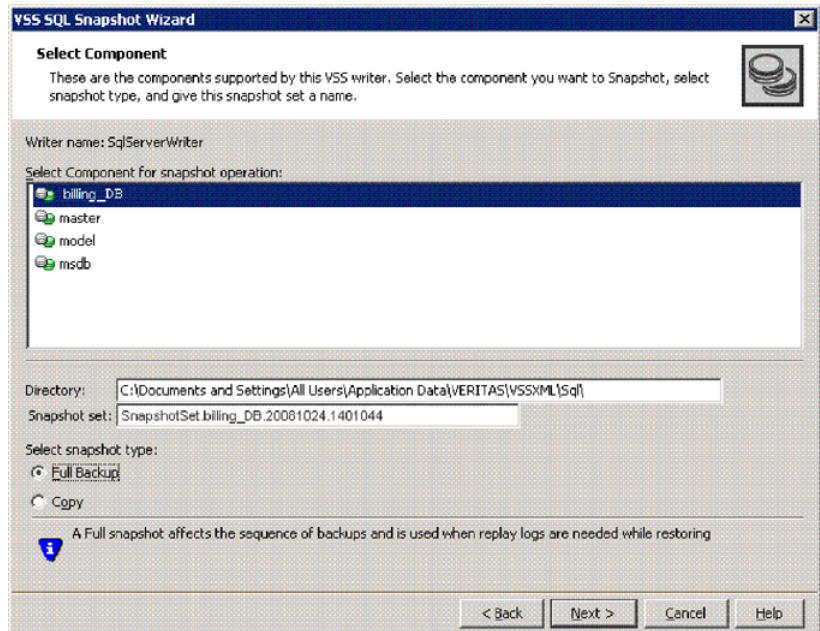
Note: To view the information about the operations that the schedule performs expand the appropriate application node to locate the schedules for that application, right-click the node, and select **Job History**.

The VSS Snapshot Scheduler wizard does not prepare the snapshot mirror. Prepare the snapshot mirror with the **Prepare** command before running the VSS Snapshot Scheduler wizard.

Note: The Prepare command replaces the Snap Start command in the GUI.

To schedule a snapshot for a selected component

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
If the Applications node is not shown, start the SQL Server VSS Writer service and then select the Storage Agent node and refresh the VEA display (**Actions > Refresh**).
- 3 Expand the **SQL** node.
- 4 Right-click the instance, and click **Schedule VSS Snapshot**.
- 5 In the **Welcome** panel, review the information and click **Next**.

6 Specify the snapshot set parameters as follows and then click **Next**:

Select component for snapshot operation Select the database for the snapshot set.

Directory

The XML file is stored by default in the directory that is shown on the screen.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`. **Save the `redirect.txt` file in the default directory, for example: `C:\Program Files\Veritas\Veritas Volume Manager\VSSXML`.**

Snapshot set	<p>Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, date, and time.</p> <p>The wizard creates the snapshot set metadata XML file with this name, with the prefix "VM_".</p>
Select snapshot type	<p>Select the snapshot type.</p> <p>You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay the logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.</p>

7 In the **Change Attributes** panel, optionally change the attributes for the snapshot volumes and click Next:

Volume Name	Displays the name of the volume.
Snapshot Volume Label	Displays the read-only label for the snapshot volume.
Drive Letter	<p>Optionally, click a drive letter and select a new choice from the drop-down menu.</p> <p>If the specified drive letters are not available when the snapshot is taken, the snapshot operation is performed, but no drive letters are assigned.</p>
Plex	Optionally, click a plex and select a new choice from the drop-down menu.

- 8 Optionally, in the **Synchronized Snapshot** panel (Volume Replicator only), select the secondary hosts for which you want to create synchronized snapshots. Either double-click on the host name or click the **Add** option to move the host into the Selected Secondary Hosts pane. To select all the available hosts, click the **Add All** option. The VSS wizard creates synchronized snapshots on all the selected secondary hosts.

This panel is displayed only in an environment using Volume Replicator (Volume Replicator). Otherwise, you are directly taken to the **Schedule Information** panel.

See *Volume Replicator Administrator's Guide*.

- 9 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

The screenshot shows the 'VSS SQL Snapshot Scheduler Wizard' window, specifically the 'Schedule information' panel. The 'General Options' tab is selected. The window title is 'VSS SQL Snapshot Scheduler Wizard'. Below the title bar, there is a sub-header 'Schedule information' and a note: 'Specify Schedule information for snapshots. Times entered on the general tab are times for the host.' There are four tabs: 'General Options', 'Days of Week', 'Days of Month', and 'Specific Dates'. The 'General Options' tab contains the following fields and options:

- Name of this schedule:
- Description of this schedule:
- Start time: End time:
- Schedule takes effect on:
- Restart task every minutes on each run day.
- Every days. Start on:
- Pre Command: Full path of command to run before snapshot.
- Post Command: Full path of command to run after snapshot.

At the bottom of the window, there are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Name of this schedule	Enter a unique name for the snapshot set schedule. This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the VSS writer name and a numbered suffix that increments with each schedule.
Description of this schedule	Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.
Start Time	The time of the day to begin taking snapshots.
End Time	The time of day to end taking snapshots. If a snapshot is in progress it is completed but a new one is not started after the end time.
Schedule takes effect on	The date on which the specified schedule takes effect. The default is the current date.

Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

10 To specify run days for the schedule, make selections on the following tabs:

Days of Week	<p>Select one or more days on one or more weeks of the month.</p> <p>You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.</p>
Days of Month	<p>Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.</p>
Specific Dates	<p>Select one or more specific dates to include in or to exclude from the schedule.</p> <p>Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.</p>

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

- 11 Click **Next**.
- 12 Review the snapshot set and schedule details and click **Finish**.

Displaying the status of the scheduled VSS SQL snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot.
The scheduled snapshots are listed in the pane on the right.
- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click the selected application and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the dialog box, view the schedule information.
You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a VSS SQL snapshot

The following is the procedure to delete a schedule.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the GUI console.

Note: The VEA GUI can delete snapshot schedules only; it does not delete mirror preparation that is scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA GUI does not update the template settings that are created with the Quick Recovery Configuration Wizard.

To delete a schedule with the GUI

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot.
The scheduled snapshots are listed in the pane on the right.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Using the VSS Restore Wizard

SFW provides recovery support for a SQL database. Using the VSS Restore Wizard, the snapshots that are created with the VSS Snapshot wizard can be used to restore the database with or without logs.

Restoring the database with logs (recovery and log replay) restores the database and the transaction log volumes. Backup transaction logs are also applied. This leaves the database in an online state.

Restoring the database without logs (recovery and no log replay) restores the database and the transaction log volumes. No other logs are applied. This leaves the database in an online state.

Restoring the database with no recovery (no recovery and no log replay) restores the database and the transaction log volumes. However no logs are applied. This leaves the database in a loading state and ready for backup transaction logs to be applied.

Restoring the database with recovery and logs

The following is the procedure to restore the database with recovery and logs.

To use log replay for an automatic roll-forward recovery to a point of failure

- 1 Ensure that you have backed up the transaction logs within SQL Server using the "overwrite existing media" option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **Applications** node.

- 5 Right-click **SQL** and click **VSS SQL Restore**.
- 6 Review the Welcome page and click **Next**.
- 7 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.
- 8 On the **Select Restore Type** panel, do the following and click Next:
 - Click **Recovery + Log replay**.
 - Enter the full path of each log file, in the order they are to be applied.
- 9 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 10 Verify the restore specifications and click **Finish**.

After the most recent backup log is replayed, the SQL Server database is closed and left in an operational state. If you took it offline earlier, bring it back online.
- 11 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS Snapshot Wizard to create a new snapshot of all the volumes in the database.

Restoring the database with recovery without logs

The following is the procedure to restore the database with recovery without logs.

To recover a database without log replay

- 1 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click **SQL** and click **VSS SQL Restore**.
- 5 Review the Welcome page and click **Next**.

- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the **Select Restore Type** panel, click **Recovery**.
- 8 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.

The database is restored to the time the snapshot set was created or last refreshed. If you took it offline earlier, bring it back online.

- 10 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS Snapshot Wizard to create a new snapshot of all the volumes in the database.

Restoring the database with one or more missing volumes

The following is the procedure to restore the database with a missing volume.

To recover a database with one or more missing volumes

- 1 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the volumes. It is also recommended to bring the database offline.
- 2 From the VEA console, navigate to the system where the database volumes are located.
- 3 Expand the system node, the Storage Agent node, and the **Applications** node.
- 4 Right-click **SQL** and click **VSS SQL Restore**.
- 5 Review the Welcome page and click **Next**.
- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all the information that is required to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the **Select Restore Type** panel, click **Recovery** and select **Recovery with missing original volume**.

Note: COW snapshots are not supported for recovery with missing volumes.

- 8 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.

The snapshot of the missing volume is changed from a read-only volume to a read-write volume.
- 10 If you have not already done so, ensure that the drive letter or mount path of the missing production volume is assigned to the snapshot volume.
- 11 Bring the database online.

If the production volume was missing, the snapshot volume is now changed to the production volume. The database is restored to the time the snapshot set was created or last refreshed.
- 12 To ensure that another split-mirror snapshot set is immediately available, use the VSS Snapshot Wizard to create a new snapshot of all the volumes in the database.

Restoring the database with no recovery

The following is the procedure to restore the database with the No Recovery option.

To restore using the No Recovery option

- 1 Ensure that you have backed up the transaction logs within SQL Server using the "overwrite existing media" option to create uniquely-named backup files.
- 2 Close the SQL GUI and all Explorer windows, applications, consoles, or third-party system management tools that may access the database volumes. It is also recommended to bring the database offline.
- 3 From the VEA console, navigate to the system where the database volumes are located.
- 4 Expand the system node, the Storage Agent node, and the **Applications** node.
- 5 Right-click **SQL** and click **VSS SQL Restore**.

Review the Welcome page and click **Next**.

- 6 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all information that is required to restore the snapshot set, including the names of the database and transaction logs volumes. Click the appropriate header to sort the list of available files by **File Name** or **Creation Time**.

- 7 On the **Select Restore Type** panel, click **No Recovery** and click **Next**.
- 8 You may receive a message "Some volumes in this component have open handles. Do you want to override these handles and do this restore? Click **Yes** to proceed." Click **No**, close any open handles and retry the command.
- 9 Verify the restore specifications and click **Finish**.

The database and log snapshot volumes are restored and the SQL Server database is left in a loading state.

- 10 Use your preferred method to manually restore the backup transaction logs to the desired point in time and then bring back the database online.
- 11 The restore operation leaves the snapshot volumes snapped back to the production volumes. To ensure that another split-mirror snapshot set is immediately available, use the VSS Snapshot Wizard to create a new snapshot of all the volumes in the database.

Copy on Write (COW)

This section describes how Storage Foundation works with Microsoft Copy on Write (COW) to provide snapshots.

COW overview

SFW supports Microsoft Copy on Write (COW) snapshots. The VEA GUI or the vxsnap CLI utility can be used to take COW snapshots. Microsoft COW snapshots are faster than full snapshots and take up less storage space.

A COW snapshot creates a shadow copy that is differential. When a change to the original volume occurs, the block that is about to be modified is read and then written to a "differences area", which preserves a copy of the data block before it is overwritten with the change. This approach stores a copy of the block before it is changed. Using the blocks in the differences area and unchanged blocks in the original volume, a copy of the volume can be constructed that represents the volume in the state before any changes occurred.

The benefit of the COW snapshot is that it creates shadow copies very rapidly because it writes only the changed blocks at a given time.

SFW VSS COW snapshot process

The process for creating COW snapshots with VSS is a two-step process. The first step is to create a Shadow Storage area (differences area). A Shadow Storage area must exist before taking a COW snapshot to preserve data blocks before they are overwritten. The second step is to create the COW snapshot.

The process for creating COW snapshots with VSS can be done with the GUI or the command line.

For more information about the command line operations, see the vxsnap CLI command.

Managing Shadow Storage

With the Shadow Storage dialog, you may add, remove, or specify another size for the Shadow Storage area.

- Add Shadow Storage lets you add a shadow storage area for a volume.

Note: Shadow storage areas are supported only on NTFS and ReFS volumes. Please note that VSS COW snapshot is not supported if the Diff area or Shadow Storage of an NTFS volume is on an ReFS volume.

- Remove Shadow Storage lets you remove the shadow copy storage area for a volume.
- Resize Shadow Storage lets you update the shadow copy storage area maximum size for a volume. This dialog lets you specify an unlimited or exact volume size.

To add Shadow Storage

- 1 Right-click the volume you want to add shadow storage.
- 2 Click **Shadow Storage > Add** in the context menu.
- 3 Select the volume to store shadow copies and the disk space size.
 - Choose the No limit option to specify that unlimited disk space can be used to store shadow copies.
 - Choose the Use limit option to set the maximum amount of disk space that is allowed to store shadow copies. Also specify the exact amount of disk space (MB) to be used.
- 4 Click **OK** to complete the settings.

Note: The `vxsnap diffarea` command can also be used to create the Shadow Storage area.

Note: A basic volume cannot be used as a Shadow Storage area.

To resize Shadow Storage

- 1 Right-click the volume that contains the shadow storage you want to resize.
- 2 Click **Shadow Storage > Resize** in the context menu.
- 3 Choose the disk space size.
 - Choose the No limit option to specify that unlimited disk space can be used to store shadow copies.
 - Choose the Use limit option to set the maximum amount of disk space that is allowed to store shadow copies. Also specify the exact amount of disk space (MB) to be used.
- 4 Click **OK** to complete the command.

To remove Shadow Storage

- 1 Right-click the volume that contains the shadow storage you want to remove.
- 2 Click **Shadow Storage > Remove** in the context menu.
- 3 Click **Yes** to complete the command.

Using the VSS COW Snapshot wizard

Perform a COW snapshot with the **COW Snap Shot** wizard.

To create VSS COW snapshots

- 1 Right-click a volume in the tree-view.
- 2 Click **COW > Snap Shot** in the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.

- 4 In the **Select Volumes** window that appears, select the volumes that you want to take a snapshot of.

Highlight the volumes in the **Available volumes** pane and click **Add**. Your selected volumes should appear in the **Selected volumes** pane.

The list of available volumes is a list of all volumes that have a shadow storage area and are in the same disk group of the selected volume.

Specify the snapshot set parameters as appropriate and then click **Next**.

Directory	The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory that is shown on the screen.
Snapshot set	Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet.cow.ID_number", date, and time.

- 5 Review the specifications of the snapshot set and click **Finish**.

Note: You can also use the `vxsnap create` CLI command to take the snapshot.

Using the VSS COW Snapshot Scheduler wizard

You can use the VSS COW Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of taking snapshots. At the scheduled time, the specified volumes are automatically snapshotted. A new snapshot set is made every time the schedule calls for a snapshot to be taken. In this way, a series of snapshot sets can be made with the scheduler. Once configured and applied, the schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

Note: You can view the information about the operations that the schedule performs by expanding the nodes under the Scheduled Tasks node in the tree-view of the GUI. Expand the Generic node to locate the schedules for snapshots of volumes. Right-click the node and select **Job History**.

Note: You must create a shadow storage area before running the VSS COW Snapshot Scheduler wizard.

To create a schedule for VSS COW snapshots

- 1 Right-click a volume in the tree-view.
- 2 Click **COW> Schedule COW SnapShot** in the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.
- 4 In the **Select Volumes** window that appears, select the volumes that you want to take a snapshot of.

Highlight the volumes in the **Available volumes** pane and click **Add**. Your selected volumes should appear in the **Selected volumes** pane.

The list of available volumes is a list of all volumes that have a shadow storage area and are in the same disk group of the selected volume.

Specify the snapshot set parameters as appropriate and then click **Next**

Directory	The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory shown on the screen.
Snapshot set	Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet.cow" date, and time. Every time a snapshot is taken, a unique snapshot set name is generated.

5 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

Name of this schedule	<p>Enter a unique name for the snapshot set schedule.</p> <p>This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of a prefix, "Generic", and a numbered suffix that increments with each schedule.</p>
Description of this schedule	<p>Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.</p>
Start Time	<p>The time of the day to begin taking snapshots.</p>
End Time	<p>The time of day to end taking snapshots.</p> <p>If a snapshot is in progress it is completed but a new one is not started after the end time.</p>
Schedule takes effect on	<p>The date on which the specified schedule takes effect. The default is the current date.</p>
Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

6 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

7 Click **Next**.

8 Review the specifications of the snapshot set and click **Finish**.

Displaying the status of the scheduled VSS COW snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1** From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2** Expand the system node and the Storage Agent node.
- 3** Click **Scheduled Tasks**.

- 4 Expand the Generic node. The scheduled snapshots are listed in the pane on the right.

All scheduled snapshots that are not associated with an Application are displayed under the Generic node.

- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click **Generic** and click **All Job History**
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the dialog box, view the schedule information.

You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a VSS COW snapshot

The following is the procedure to delete a schedule.

You can also delete (but not modify) a schedule from the VEA GUI console.

To delete a schedule with the VEA GUI

- 1 From the VEA console, navigate to the system where the production volumes are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks**.
- 4 Expand the Generic node for which you scheduled the snapshot. The scheduled snapshots are listed in the pane on the right.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Using the VSS COW Restore Wizard

Uses the snapshot volume in a snapshot set created s by VSS COW snapshot to restore data, for example, after an original volume has become corrupted.

To restore a volume with COW snapshot

- 1 Right-click a volume or a component of an application in the tree-view.
- 2 Click **COW > Restore** in the context menu.
- 3 In the wizard, review the Welcome page and click **Next**.

- 4 Select the snapshot set XML metadata file to be used for this operation and click **Next**.

The XML metadata file contains all the information that is required to restore the volume.
- 5 In the **Select Volumes** window that appears, select the volumes that you want to restore.

Highlight the volumes in the **Available volumes** pane and click **Add**. Your selected volumes should appear in the **Selected volumes** pane.
- 6 Review the specifications of the snapshot set and click **Finish**.

Additional information about COW snapshots

Since COW snapshots rely on a Shadow Storage area, COW snapshots have the following constraints:

- COW snapshots are developed by keeping track of the blocks that have changed. If the original volume that the COW snapshot is based upon becomes unavailable or corrupted, then the COW snapshot cannot be used to restore the snapshotted volume.
- After restoring the original volume with a COW snapshot, any subsequent COW snapshots that were taken after that snapshot are automatically deleted by Windows.
- The size of the allocated shadow storage area limits the number of COW snapshots for a volume. When the shadow storage area has reached its capacity, then older snapshots are deleted to make room for new snapshots. In addition, 512 COW snapshots is the maximum number allowed per volume.
- During a restore operation, a forced dismount of the original volume or the volume containing the shadow storage area may result in I/O corruption and the loss of COW snapshots.
- Due to the incremental nature of COW snapshots, COW snapshots can only be used to restore the snapshotted volume to a point in time.
- For every write operation on a new block of the snapshotted volume, a read operation is performed to allow the COW snapshot to store the block in the shadow storage area. As a result, a COW snapshot may affect the I/O performance of the snapshotted volume.
- SFW only supports COW snapshots of dynamic volumes.
- Both the shadow storage area and the snapshotted volume must reside in the same dynamic disk group.

- If a volume containing a shadow storage area is a part of a snapshot set and the shadow storage area volume is restored, then all snapshots that were stored on that volume are deleted automatically by Windows. Make sure that volumes containing shadow storage areas for other volumes are not part of the same snapshot set. Otherwise COW snapshots may automatically get deleted during a restore operation, resulting in a snapshot not found error.

After performing operations for COW snapshots or operations on the Shadow Storage area, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

Using the VSS COW snapshot wizards with Microsoft Exchange

FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all the volumes that are associated with an Exchange mailbox database without taking the database offline. The VSS COW Snapshot wizard uses VSS to take snapshots of mailbox databases. The VSS process allows the storage group's databases to be quiesced before the snapshot operation and reactivates them immediately after.

FSW also provides a VSS COW Snapshot Scheduler wizard that can be used as an alternative to the VSS COW Snapshot wizard. It enables you to set up a schedule for taking the snapshots at the specified times. The schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

Using the VSS COW Snapshot wizard

Creating a snapshot is a two-step process. The first step, creating a shadow storage area for the snapshot, creates a differences area to store blocks of data before they are changed. The second step uses the VSS COW Snapshot wizard to take the snapshot and create an XML file to store the Exchange and snapshot volume metadata. The VSS COW Snapshot wizard integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes in the storage group. This snapshot is done while the databases are online and without disrupting the email flow.

Note: You must create a shadow storage area before running the VSS COW Snapshot wizard.

Note: After performing COW-related operations with the Windows GUI, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

Creating the snapshot set

The following is the procedure to create a snapshot set.

Note: You must create a shadow storage area before running the VSS COW Snapshot wizard.

To create the snapshot set using the GUI

- 1 From the VEA console, navigate to the system where the production volumes are located.
- 2 Expand the system node, the Storage Agent node, and the **Applications** node.
- 3 Right-click Exchange and click **VSS COW Snapshot**
- 4 In the wizard, review the Welcome page and click **Next**.
- 5 Specify the snapshot set parameters as follows and then click **Next**

Select Component for snapshot operation	Select the component for the snapshot set. For Exchange, the component is a storage group.
Directory	The snapshot set metadata XML file is stored under this name. The XML file is stored by default in the directory that is shown on the screen. However, in a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster. To accomplish this, use a text editor to create a text file named "redirect.txt." This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the metadata file, for example, <code>\\ServerName\SharedFolderName. Save the redirect.txt file in the default directory C:\Program Files\Veritas\Veritas Volume Manager\VSSXML on each node of the cluster.</code>
Snapshot set	Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, identification number, date, and time.

Select snapshot type Select the snapshot type.

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the selected component
- Runs **Eseutil** to check for consistency before truncating the logs
- Truncates the transaction logs

Copy is typically used for Quick Recovery. It creates a copy of the storage group, but does not truncate the transaction logs. Optionally check **Run Eseutil** with the **Copy** option to check the snapshot for consistency.

6 Review the specifications of the snapshot set and click **Finish**

Using the VSS COW Snapshot Scheduler wizard

You can use the VSS COW Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of taking snapshots. At the scheduled time, the specified volumes are snapshotted automatically. Once configured and applied, the schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background. In a clustered server environment, ensure that the scheduler service is configured on each node with domain administrator privileges so that any node in the cluster can run the schedule after a failover.

Note: You can find information about the operations that the schedule performs by expanding the nodes under the Scheduled Tasks node in the tree-view of the GUI. Expand the appropriate application node to locate the schedules for that application. Right-click the node and select **Job History**.

Note: You must create a shadow storage area before running the VSS COW Snapshot wizard.

To schedule a snapshot for a selected component

- 1** From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2** Expand the system node, the Storage Agent node, and the **Applications** node.
- 3** Right-click the Exchange storage group, and click **Schedule VSS COW Snapshot**.

- 4 In the **Welcome** panel, review the information and click **Next**.
- 5 Specify the snapshot set parameters as follows and then click **Next**:

Select component for snapshot operation Select the component for the snapshot set.

Directory The XML file is stored by default in the directory that is shown on the screen.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`. **Save the `redirect.txt` file in the default directory, for example: `C:\Program Files\Veritas\Veritas Volume Manager\VSSXML`.**

Note: When a COW snapshot is scheduled, a snapshot is taken at the specified time and a new xml file is generated for each snapshot.

Snapshot set Enter a name for the snapshot set, for example, billing or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, identification number, date, and time.

Select snapshot type Select the snapshot type.
You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database.

6 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

Name of this schedule	<p>Enter a unique name for the snapshot set schedule.</p> <p>This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the application name and a numbered suffix that increments with each schedule.</p>
Description of this schedule	<p>Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.</p>
Start Time	<p>The time of the day to begin taking snapshots.</p>
End Time	<p>The time of day to end taking snapshots.</p> <p>If a snapshot is in progress it is completed but a new one is not started after the end time.</p>
Schedule takes effect on	<p>The date on which the specified schedule takes effect. The default is the current date.</p>
Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

7 To specify run days for the schedule, make selections on the following tabs:

- | | |
|----------------|---|
| Days of Week | Select one or more days on one or more weeks of the month.

You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month. |
| Days of Month | Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month. |
| Specific Dates | Select one or more specific dates to include in or to exclude from the schedule.

Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9. |

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

8 Click **Next**.

9 Review the snapshot set and schedule details and click **Finish**.

Displaying the status of the scheduled VSS COW snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot.

The scheduled snapshots are listed in the pane on the right.

All scheduled snapshots that are associated with Exchange are displayed under the Exchange node.

- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click the selected application and click **All Job History**
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the dialog box, view the schedule information.

You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a VSS COW Snapshot

The following is the procedure to delete a schedule.

You can also delete (but not modify) a schedule from the VEA GUI.

To delete a schedule using the VEA GUI

- 1 From the VEA console, navigate to the system where the production volumes are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.

The scheduled snapshots are listed in the pane on the right.
- 4 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Using the VSS COW snapshot wizards with Microsoft SQL

SFW provides support for taking VSS COW snapshots of Microsoft SQL databases. FlashSnap integrates with the Microsoft Volume Shadow Copy Service (VSS) to allow snapshots to be taken of all volumes that are associated with an SQL database without taking the database offline. The VSS COW Snapshot wizard uses VSS to take snapshots of the database. The VSS process allows the database to be quiesced before the snapshot operation and reactivates it immediately after.

SFW also provides a VSS COW Snapshot Scheduler wizard that can be used as an alternative to the VSS COW Snapshot wizard. It enables you to set up a schedule for taking the snapshots. The schedule is maintained by a scheduler service, VxSchedService.exe, that runs in the background.

Using the VSS COW Snapshot wizard

Creating a snapshot is a two-step process. The first step, create a shadow storage area for the snapshot, creates a differences area to store blocks of data before they are changed. The second step uses the VSS Snapshot wizard to create on-host snapshot volumes as well as an XML file to store the database and snapshot volume metadata. The VSS Snapshot wizard integrates with VSS to quiesce the database and then simultaneously snapshot the volumes in the database. This snapshot is done while the database is online and without disrupting processing.

Note: You must create a shadow storage area before running the VSS COW Snapshot wizard.

Note: After performing COW-related operations with the Windows GUI, the VEA GUI may not be updated to reflect the results of the operation. Use the VSS Refresh command to update the VEA GUI.

Creating the snapshot set

The following is the procedure to create a snapshot set.

To create the snapshot set

- 1 From the VEA console, navigate to the system where the production volumes are located.
- 2 Expand the system node, the Storage Agent node, Applications node, and SQL node.
- 3 Right-click the SQL instance node and click **VSS COW Snapshot**.
- 4 In the wizard, review the Welcome page and click **Next**.

5 Specify the snapshot set parameters as follows and then click **Next**:

Select Component for snapshot operation Select the database for the snapshot set.

Directory The wizard creates the snapshot set metadata XML file. The XML metadata file is stored by default in the directory that is shown on the screen.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`. Save the `redirect.txt` file in the default directory, for example: `C:\Program Files\Veritas\Veritas Volume Manager\VSSXML`.

Snapshot set Enter a name for the snapshot set, for example, `billing` or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, identification number, date, and time.

Select snapshot type Select the snapshot type.

You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay the logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

6 Review the specifications of the snapshot set and click **Finish**.

Using the VSS COW Snapshot Scheduler wizard

You can use the VSS COW Snapshot Scheduler wizard to add a snapshot schedule.

The scheduling capability automates the process of taking snapshots. At the scheduled time, the specified volumes are automatically snapshotted. Once configured and applied, the schedule is maintained by a scheduler service, `VxSchedService.exe`, that runs in the background. In a clustered server environment, ensure that the scheduler service is configured on each node with domain

administrator privileges so that any node in the cluster can run the schedule after a failover.

Note: You can view the information about the operations that the schedule performs by expanding the nodes under the Scheduled Tasks node in the tree-view of the GUI. Expand the appropriate application node to locate the schedules for that application. Right-click the node and select **Job History**.

Note: You must create a shadow storage area before running the VSS COW Snapshot wizard.

To schedule a snapshot for a selected component

- 1** From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2** Expand the system node, the Storage Agent node, Applications node, and SQL node.
- 3** Right-click the SQL instance, and click **Schedule VSS COW Snapshot**.
- 4** In the **Welcome** panel, review the information and click **Next**.

5 Specify the snapshot set parameters as follows and then click **Next**:

Select component for snapshot operation Select the database for the snapshot set.

Directory The XML file is stored by default in the directory that is shown on the screen.

In a clustered server environment, the XML file must be saved on shared storage to be available from all nodes in the cluster.

There are two ways to change the XML file location.

First, edit the directory path in the Directory field for this wizard screen. Second, change the XML file location. Use a text editor to create a text file named `redirect.txt`. This text file should contain a single text line specifying the Universal Naming Convention (UNC) path to the location of the XML file, for example, `\\ServerName\SharedFolderName`. Save the `redirect.txt` file in the default directory, for example: `C:\Program Files\Veritas\Veritas Volume Manager\VSSXML`.

Note: When a COW snapshot is scheduled, a snapshot is taken at the specified time and a new xml file is generated for each snapshot.

Snapshot set Enter a name for the snapshot set, for example, `billing` or accept the default name. The wizard generates a default snapshot set name that includes the term "SnapshotSet", component name, identification number, date, and time.

Select snapshot type Select the snapshot type.

You can specify that snapshots be created as either a Full backup or Copy backup type. Either type can be used to restore a database. However, if you want to replay logs in SQL Server as part of restoring a database, a Full backup needs to have been created earlier. When replaying logs, you can replay from the time of the last Full backup. A Copy backup does not affect this sequence of log replay and therefore is often used as an "out of band" copy for purposes such as testing or data mining.

6 In the **Schedule Information** panel, on the **General Options** tab, specify the following:

Name of this schedule	<p>Enter a unique name for the snapshot set schedule.</p> <p>This name identifies the snapshot schedule if you later want to view information about the snapshot status. A default name consists of the application name and a numbered suffix that increments with each schedule.</p>
Description of this schedule	<p>Optionally, enter a description to help you identify the schedule when you view information about the snapshot status.</p>
Start Time	<p>The time of the day to begin taking snapshots.</p>
End Time	<p>The time of day to end taking snapshots.</p> <p>If a snapshot is in progress it is completed but a new one is not started after the end time.</p>
Schedule takes effect on	<p>The date on which the specified schedule takes effect. The default is the current date.</p>
Restart task every	<p>The interval between snapshots, in minutes.</p> <p>For example, if the interval is 360 minutes and you schedule a snapshot start time of 12:00 P.M. and an end time of 7:00 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.</p>
Every	<p>Enable the Every option to have the snapshot schedule continue to occur. Otherwise the schedule applies only for one day.</p> <p>Specify the number of days before restarting the snapshot schedule.</p> <p>For example, 1 day would mean that the schedule takes effect daily, 2 days would mean every other day.</p>
Start On	<p>If you enable the Every option, specify the starting date.</p>
Pre Command	<p>Optionally, specify the full path of a command script to run before the scheduled snapshot occurs.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p>

- 7 To specify run days for the schedule, make selections on the following tabs:

Days of Week	Select one or more days on one or more weeks of the month. You can click a button at the top of the column to select the entire column or a button to the left of a row to select the entire row. For example, clicking First schedules the snapshots to occur on the first occurrence of all the week days for the month.
Days of Month	Select one or more days of the month. You can also check the Last Day check box to schedule the snapshot for the last day of each month.
Specific Dates	Select one or more specific dates to include in or to exclude from the schedule. Excluding a date takes precedence over days scheduled on the other tabs. For example, if you schedule every Monday on the Days of Week tab, and you exclude Monday October 9 on the Specific Dates tab, the snapshots are not taken on October 9.

If two schedules overlap for the same snapshot set, only one snapshot is taken. For example, if you select every Thursday plus the last day of the month, and the last day of the month occurs on Thursday, only one snapshot is taken on Thursday.

- 8 Click **Next**.
- 9 Review the snapshot set and schedule details and click **Finish**.

Displaying the status of the scheduled VSS COW snapshot

If a scheduled snapshot fails for some reason, the scheduler process attempts to rerun it. You may want to verify that scheduled snapshots completed successfully. From the VEA console, you can view snapshot results.

To view a scheduled snapshot status

- 1 From the VEA console, navigate to the system where the production volumes and snapshot mirrors are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot.

The scheduled snapshots are listed in the pane on the right.

All scheduled snapshots that are associated with SQL are displayed under the SQL node.

- 5 Choose one of the following:
 - To view the status of all scheduled jobs, right-click the selected application and click **All Job History**.
 - To view the status of a particular schedule, right-click the snapshot schedule name and click **Job History**.
- 6 In the dialog box, view the schedule information.

You can sort listed schedules by clicking the column headings. The Status column shows if the snapshot completed successfully.

Deleting a schedule for a VSS COW snapshot

The following is the procedure to delete a schedule.

You can also delete (but not modify) a schedule from the VEA GUI.

To delete a schedule using the VEA GUI

- 1 From the VEA console, navigate to the system where the production volumes are located.
- 2 Expand the system node and the Storage Agent node.
- 3 Click **Scheduled Tasks** to view all the applications that have schedules.
- 4 Select the application for which you scheduled the snapshot.

The scheduled snapshots are listed in the pane on the right.
- 5 Right-click the name of the snapshot schedule and click **Delete Schedule**.

Configuring data caching with SmartIO

This chapter includes the following topics:

- [About SmartIO](#)
- [Typical deployment scenarios](#)
- [How SmartIO works](#)
- [SmartIO benefits](#)
- [SmartIO limitations](#)
- [About cache area](#)
- [About SmartIO caching support](#)
- [Configuring SmartIO](#)
- [Frequently asked questions about SmartIO](#)

About SmartIO

Traditional disks often pose as an I/O bottleneck for high transaction applications. You can compensate this by either increasing the in-RAM cache size or investing in expensive storage. SmartIO is a data caching solution that helps in improving I/O performance of applications by:

- Providing a read-only/write-through I/O cache on a locally attached Solid State drives (SSDs). To support VMware vMotion, the SmartDisks must replace SSDs. SmartDisks are the disks assigned by Dynamic Multi-Pathing for VMware. These disks are provisioned from an aggregated pool of SSDs.)

- Reducing the I/O requests on the back-end storage
- De-centralizing the storage
- Freeing the available network and controller capacity to handle traffic for more servers and application consolidation on a system
- Interposing a central management layer between cached-device and the applications that need to access them

SmartIO supports volume-level read-only caching as SSDs are primarily beneficial in high-read environments.

Typical deployment scenarios

SmartIO can be configured in both, clustered as well as non-clustered environments.


Note: A cache area itself cannot be configured for high availability.

The following sections depict sample cluster deployments and indicate the SmartIO behavior after a failover:

Cluster deployment in a Hyper-V environment

A cluster deployment in a Hyper-V environment involves on-host or in-guest SFW installation.

The following diagram represents on-host SFW installation. The application data is stored on the shared disks and the cache area is created on the locally attached SSD.



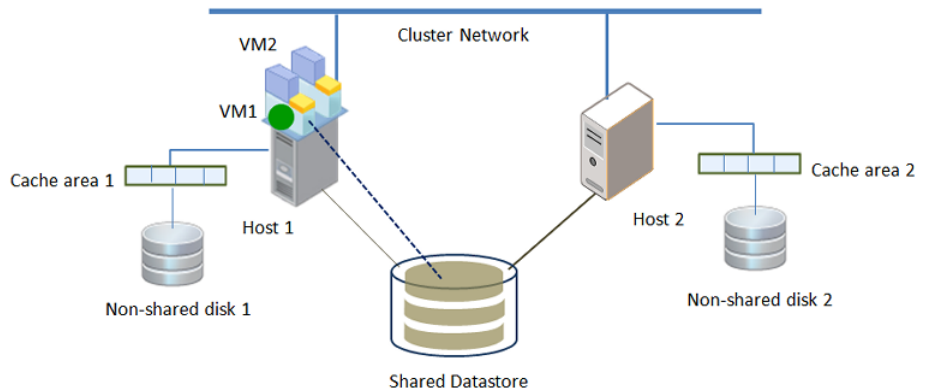
SmartIO behavior after a failover to Node 2:

- If the data volumes on Node 1 are linked to an auto cache area:
 - Links the data volumes to an auto cache area on Node 2
 - If an auto cache area is not available on Node 2, does not link the volumes to any cache area
- If the data volumes on Node 1 are linked to a noauto cache area:
 - Links the data volumes to a noauto cache area available on Node 2 (the noauto cache area with the same name as that on Node1 takes priority)
 - If a noauto cache area is not available on Node 2, link the volumes to an auto cache area that is available on Node 2

Cluster deployment in a VMware environment- No vMotion support

A cluster deployment in a VMware environment involves in-guest SFW installation.

The following diagram represents in-guest SFW installation where VMware vMotion is not supported. The application data is stored on the shared disks and the cache area is created on the locally attached SSD, that is exposed to the virtual machines.

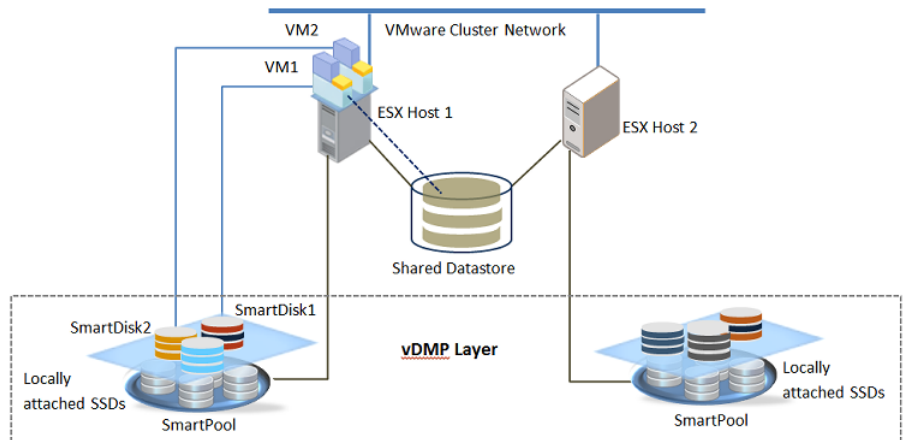


SmartIO behavior

- If the application fails over to VM2 (failover on the same host), SFW continues to serve the I/O request from the cache that was created and there is no effect on the I/O performance.
- If the VM moves to Host2 (vMotion), then the vMotion operation fails. VMware vMotion fails if any disk is locally attached to the virtual machine.

Cluster deployment in a VMware environment- with vMotion support

A cluster deployment in a VMware environment involves in-guest SFW installation. The following diagram represents in-guest SFW installation with support for VMware vMotion. The application data is stored on the shared disks and the cache area is created on the locally attached SmartDisks.



In the above example, SFW and the application that is configured for high availability is installed on the virtual machines (VM1 and VM2). A VMware cluster is configured between the ESX hosts. Multiple SSDs are locally attached to the individual ESX hosts. Using VxDMP (version 7.1 or later), these SSDs are aggregated to form individual SmartPools. From a SmartPool, multiple SmartDisks are provisioned for caching. A SmartDisk is exposed to a virtual machine as a RAW LUN (Here, SmartDisk1 is attached to VM1 and SmartDisk2 is attached to VM2).

Note: You can attach only one SmartDisk per virtual machine.

For details about creating and administering SmartPools, and provisioning the SmartDisks, refer to *Dynamic Multi-Pathing Administrator's Guide*

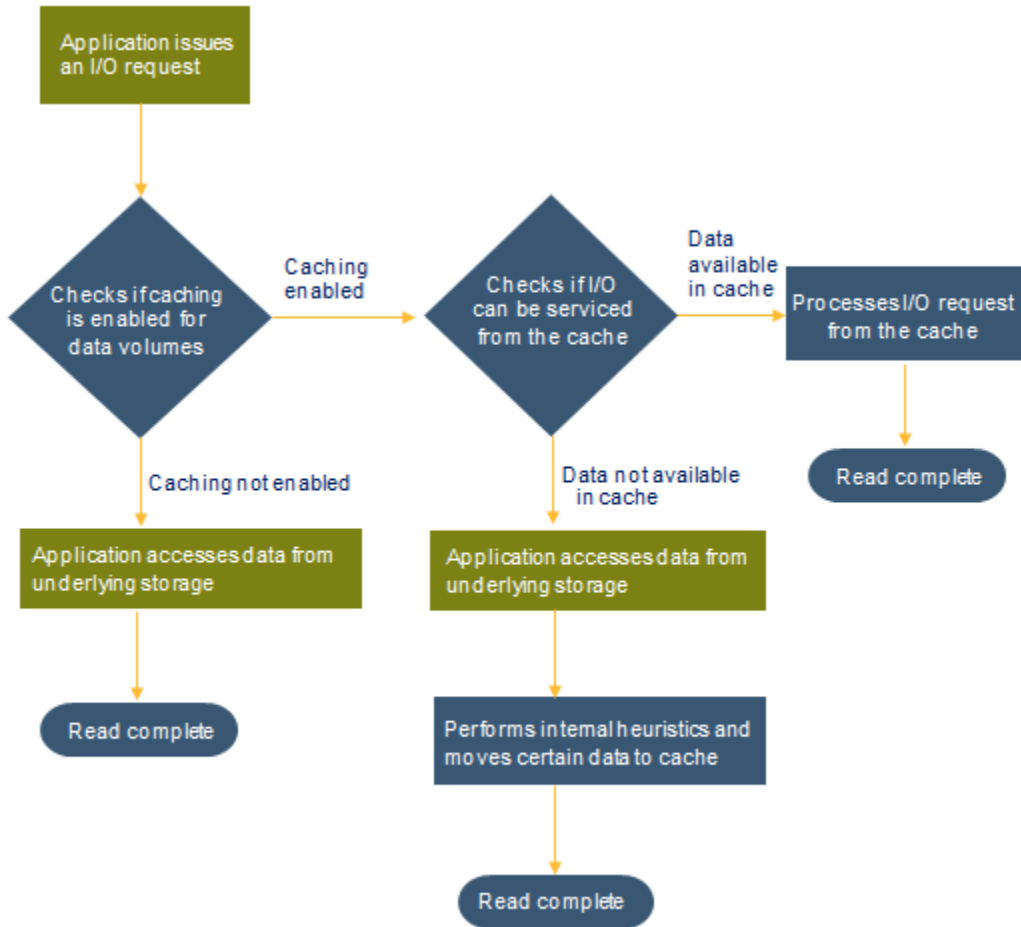
SmartIO behavior

After vMotion, the first I/O to the cache area fails and the request is forwarded to the back-end storage (data disks). SFW considers the I/O failure as a vMotion event and informs the VxDMP layer to allocate a SmartDisk of the required capacity. After the VxDMP layer allocates the SmartDisk, SFW retrieves the cache area details from the registry and creates a similar cache area on the SmartDisk. This newly created cache area is assigned to the virtual machine after vMotion.

Note: If the cache area is offline during vMotion, SFW does not recreate the cache area on the SmartDisk attached to the target host. Admin intervention is required to bring the cache area online. After the cache is brought online the subsequent I/O fails. This I/O failure is considered as a vMotion case and SFW accordingly initiates the further process.

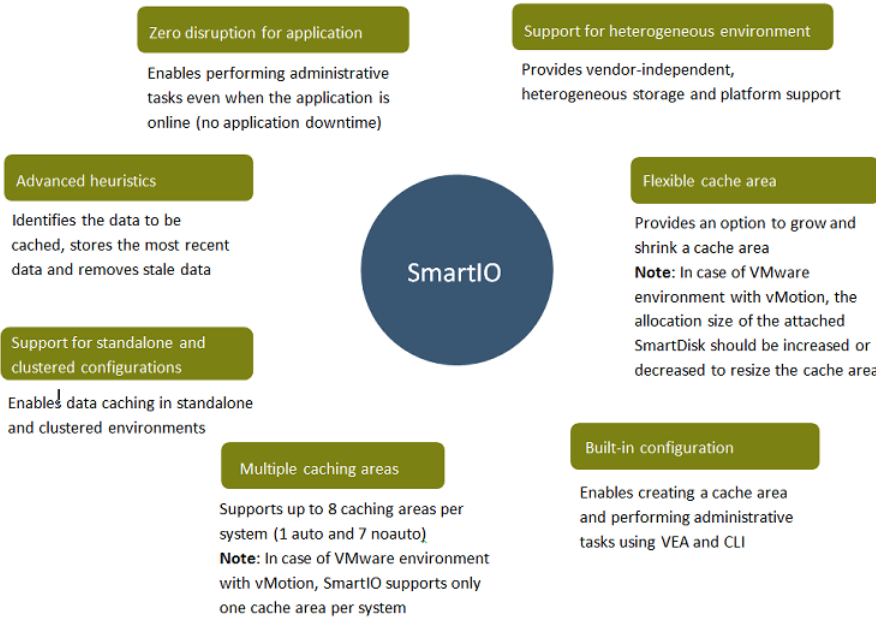
How SmartIO works

SmartIO uses a cache area on the locally attached SSD to store the cached data. The following figure represents SmartIO workflow:



SmartIO benefits

As a data caching solution, SmartIO provides the following benefits:



SmartIO limitations

SmartIO has the following limitations:

- You cannot reserve a cache area for a particular volume. You can create a new cache area and link the volume with it.
- File pinning or block pinning is not supported.
- The cache is volatile and does not persist after system restart.

About cache area

Cache area is the storage space allocated on the SSD(s) for caching. It is used to store cache data corresponding to any caching-enabled volume. If you use an SSD for caching, you cannot use it for storage.

Attributes of the cache area

Support for automatic caching is a characteristic of the cache area. The attributes of a cache area define its association type.

The attributes of a cache area are:

auto

This attribute sets the cache area for automatic caching. All volumes on the system are cached unless you explicitly disable caching for that volume. Storage Foundation (SFW) supports a single auto cache area. You can choose not to have an auto cache area for your system.

noauto

The cache area is not enabled for automatic caching. You must explicitly link the volumes to the cache area with noauto attribute.

About SmartIO caching support

The following table lists the SmartIO conditions that SmartIO supports and the ones that it does not support:

Table 12-1 Conditions SmartIO supports and does not support

Supports	Does not support
All the existing volumes which are cache-enabled and are not linked to any specific cache area	RAID-5 and log volumes (such as DCO, SRL, and DCM)
All the new volumes that are created on the primary node	File pinning or block pinning
Volumes that fail over from primary node to secondary node in a cluster environment	Reserving a cache area for a particular volume
Volumes of upgraded disk group when disk group version is upgraded to 7.1	More than 8 cache areas per system Note: In case of VMware environment with vMotion, only one cache area is supported per system.
Any SSD device that VxVM supports (must be initialized for use with VxVM)	Cache area that is created on shared SSD
Disk group version 171 or above (for volumes to be cached)	High availability of a cache area
Volumes that are created on secondary or clustered disk group and fast-Failover enabled disk group on active node	Multiple cache areas on single SSD

Table 12-1 Conditions SmartIO supports and does not support (*continued*)

Supports	Does not support
Online shared volumes in clustered disk group, snapshot volumes, and Volume Replicator volumes on primary site	Persistent cache data (cache data available even after a system power cycle)
Caching after vMotion	

Configuring SmartIO

SmartIO is enabled by default when you install InfoScale Storage or InfoScale Enterprise. To use SmartIO, you create a cache area using one or more non-shared SSDs, and link the VxVM volumes to the cache area. In case of VMware environment with vMotion, SmartDisks replace the SSDs.

Notes:

- The size of the SSDs or the SmartDisk used for caching defines the cache area size. The SSDs or the SmartDisks that are used for caching are added in a CACHEPOOL and cannot be used as storage disks.
- All the data volumes that are created are implicitly associated with the "auto" cache area. You can change the association to "noauto" and then explicitly associate the data volumes to the cache area. In case of multiple cache areas, only one cache area has an auto association.
- Veritas recommends that the SmartDisks must not be used as data disks.
- As part of the earlier configuration, if you have cache areas created on a locally attached SSD and plan to use SmartIO in VMware vMotion environment, you must delete the existing cache areas and create them on the SmartDisk attached.
- If the data volume is disconnected, caching for that volume is stopped. Caching is restarted afresh once the volume is reconnected and brought online. If the cache area is disconnected, the cache area is taken offline and hence it stops caching of all the volumes that are linked with it.
- If vMotion is initiated when the cache area is offline, SFW does not re-create the cache area. Admin intervention is required to reconnect the disconnected SmartDisk and to bring the cache area online. SFW then performs the operations that are required for recreating the cache area.
- You cannot take or restore a snapshot of the virtual machine to which a SmartDisk is attached. Before taking or restoring the snapshot, you must delete the cache area and unassign the SmartDisk, using the vDMP client. After the snapshot is taken or

restored, you must reassign the disk using the vDMP client and then recreate a cache area.

You can create a cache area and perform the administrative tasks using VEA and CLI.

Creating and administering cache area using VEA

The following SmartIO tasks can be performed using VEA:

See [“Creating a cache area”](#) on page 391.

See [“Enabling caching on a volume”](#) on page 392.

See [“Growing a cache area”](#) on page 392.

See [“Shrinking a cache area”](#) on page 393.

See [“Disabling caching on a volume”](#) on page 393.

See [“Deleting a cache area”](#) on page 393.

See [“Viewing cache area properties”](#) on page 394.

Creating a cache area

You can create a cache area on the locally attached SSD using the Create Cache Area Wizard.

You can use internal or external SSDs, or a combination of both to create the cache area. You can create maximum 8 cache areas and can link it to a specific set of volumes.

To create a cache area

- 1 Right-click Cache in the left panel tree view and select **Create Cache Area** from the context menu.
- 2 The Create Cache Area wizard appears. On the **Welcome** panel, click **Next**.

- 3 On the **SSD Selection** panel, enter a name in the **Cache Area Name** box. A cache area name is limited to 15 ASCII characters. It cannot contain spaces, or forward or backward slashes. Also, a period cannot be the first character.

The SSDs attached to the system are displayed in the **Available disks** list. Select the SSDs to be used in the cache area.

To set the cache area as the auto cache area, select the **Enable default caching** check box. Click **Next**.

Note: SmartIO supports only one auto cache area. To set a new default cache area when a default cache area is already set on your system, you must reset the existing default cache area through CLI.

- 4 On the **Summary** panel, review the summary report and click **Finish**.

Enabling caching on a volume

You can enable caching for a volume and link it implicitly with the auto cache area or specify a noauto cache area.

To enable caching on a volume

- 1 Expand the Volumes node in the tree view. Right-click the desired volume and select **Enable Caching** from the context menu.
- 2 In the dialog box that appears, select the cache area to be linked with the volume. To link the volume implicitly with the auto cache area, select the **Use default cache area** check box. Click **OK**.

Growing a cache area

You can grow a cache area by adding SSDs to it.

Note: You cannot perform this operation in case of VMware environment with vMotion. To grow a cache area in VMware environment, you must delete the created cache area, and then using the vDMP client increase the size of the SmartDisk attached. After you increase the SmartDisk size, you must create a cache area with the desired size.

To grow a cache area

- 1 Expand the Cache node in the tree view. Right-click the cache area to be grown and select **Grow Cache Area** from the context menu.
- 2 The Grow Cache Area wizard appears. On the **Welcome** panel, click **Next**.

- 3 On the **SSD Selection** panel, select the cache area to be grown. The SSDs that can be added to the cache area are displayed in the **Available disks** list. Select the SSDs to be added to the cache area. Click **Next**.
- 4 On the **Summary** panel, review the summary report and click **Finish**.

Shrinking a cache area

You can shrink a cache area by removing SSDs from it.

Note: You cannot perform this operation in case of VMware environment with vMotion. To shrink a cache area in this environment, you must delete the created cache area, and then using the vDMP client decrease the size of the SmartDisk attached. After you decrease the SmartDisk size, you must create a cache area with the desired size.

To shrink a cache area

- 1 Expand the Cache node in the tree view. Right-click the cache area to be shrunk and select **Shrink Cache Area** from the context menu.
- 2 The Shrink Cache Area wizard appears. On the **Welcome** panel, click **Next**.
- 3 On the **SSD Selection** panel, select the cache area to be shrunk. The SSDs that can be removed from the cache area are displayed in the **Available disks** list. Select the SSDs to be removed from the cache area. Click **Next**.
- 4 On the **Summary** panel, review the summary report and click **Finish**.

Disabling caching on a volume

When you disable caching on a volume, the volume is unlinked from the cache area.

If the cache-disabled volume fails over to any other cluster node, it remains cache-disabled.

To disable caching on a volume

- 1 Expand the Volumes node in the tree view. Right-click the volume and select **Disable Caching** from the context menu.
- 2 In the dialog box that appears, click **Yes**.

Deleting a cache area

Deleting a cache area unlinks all volumes that are linked with the cache area. If another cache area with noauto attribute is present, then these volumes are linked

with that cache area. If another cache area with noauto attribute is not present, and an auto cache area is present, then these volumes are linked with the auto cache area.

To delete a cache area

- 1 Expand the Cache node in the tree view. Right-click the cache area to be deleted and select **Delete Cache Area** from the context menu.
- 2 In the dialog box that appears, click **Yes**.

Viewing cache area properties

You can view the properties of a cache area by right-clicking the cache area and selecting Properties from the context menu. The Properties window comes up.

The categories for the cache area that are shown in the sample screen are described in the table that follows:

Table 12-2 Cache area properties

Category	Description
Cache area name	The name of the cache area
State	Whether the cache area is Online or Offline
Size	Size of the cache area
Default	Whether the cache area attribute is auto or noauto

Creating and administering cache area through CLI

The following SmartIO tasks can be performed using the command line options:

- Create a cache area.
See [“sfcache create”](#) on page 751.
- Delete a cache area.
See [“sfcache delete”](#) on page 751.
- Grow a cache area.
See [“sfcache growby”](#) on page 753.
- Shrink a cache area.
See [“sfcache shrinkby”](#) on page 753.
- Set or reset a cache area as auto cache area.
See [“sfcache set”](#) on page 753.

- Bring cache area offline.
You can bring a cache area offline for maintenance purposes, for instance, for repairing frequent connect or disconnect of the SSD that the cache area uses. When a cache area is offline, the caching for all the volumes which are linked with the cache area is stopped, but the linking remains active. After the cache area is brought online, caching is restarted for the linked volumes.
If you bring the auto cache area offline, the caching is stopped for all the linked volumes. After the auto cache area is brought online, caching is restarted on the linked volumes, as well as for the volumes that are caching-enabled.
See “[sfcache offline](#)” on page 752.
- Bring cache area online.
You can bring a cache area online to restart the caching for the volumes that are linked with it. The stale cache is invalidated for the volumes that are linked with the cache area before restarting caching.
See “[sfcache online](#)” on page 752.
- Enable caching on a volume
See “[sfcache enable](#)” on page 751.
- Disable caching on a volume
See “[sfcache disable](#)” on page 752.
- List all the cache areas present on the system
See “[sfcache list](#)” on page 753.
- Display the statistics of a cache area
See “[sfcache stat](#)” on page 754.

Frequently asked questions about SmartIO

This section lists the frequently asked questions about SmartIO and provides detailed answers on each of these questions.

How to configure a volume to use a non-default cache area?

By default, a data volume gets linked to the auto cache area. To configure a volume to use a non-default cache area, create a noauto cache area on the SSD attached to the system and disable caching for the appropriate data volume. After you create a noauto cache area, enable caching and explicitly link the volume to it.

What is write-through I/O caching?

Write-through caching and read caching is synonymous. In both these caching types, all the reads are cached. The write data is written to cache only if it is present in the cache.

Does SmartIO with SFW support write caching?

If I/O-specific data is not present in a cache, it is written to the data volumes. SmartIO performs internal heuristics to decide if the data needs to be cached as well. SmartIO with SFW does not support writing (write back) data directly to the cache.

Are there any logs that I can refer to, if caching fails for a particular volume?

If caching fails for any volume, you can refer to the Vold and VxIO log files for details. The log files are available at the following location on the system:

```
%vmpath%\vxio.etl %vmpath%\logs\voldlog.txt
```

where, %vmpath% is typically C:\Program Files\Veritas\Veritas Volume Manager\Logs

I have deleted a cache area, but the disk is still present in the Cachepool. How can I remove it from the Cachepool?

When you create a cache area, the SSD is placed in the cache pool. If the operation fails, then SmartIO should remove the SSD from the cache pool to allow the user to retry creating a cache area with the SSD. But if this issue occurs, then the SSD is not removed properly from the cache pool, and hence it cannot be added to the cache area. Similarly, if a cache area is deleted, SmartIO should remove the SSD in the cache area from the cache pool. Now if the operation fails, the SSD is not removed successfully from the cache pool.

To remove the disk from the Cachepool, run the following command:

```
vxdbg -gcachepool -f rmdisk <diskname>
```

Is the VxVM cached data persistent?

No, the cached data is not persistent. It is deleted after a system power-cycle.

Is an application's performance affected if the cache device becomes inaccessible while caching is enabled?

If a cache device becomes inaccessible while caching is enabled, the application continues to function normally. However, the I/O performance may be reduced. Until the device is reconnected, the I/Os are served from the data disks.

Are there any tools available to measure SmartIO performance?

You can use tools such as `perfmon` or `typeperf` to measure SmartIO performance based on the number of reads received. For statistics on the cache area, you can use the `sfcache stat` tool.

Will there be a performance drop after vMotion?

Yes, after vMotion there is a slight drop in I/O performance. A cache area is created on the SmartDisk that is attached to the target host, after the first I/O that is issued post vMotion fails to fetch the data from a cache. This I/O is served from the data disks. The subsequent I/Os are served from the data disks until that data is stored in the cache.

Will the cache area recreation fail, if the SmartDisk assigned has insufficient space?

Yes, the cache area recreation fails, if the SmartDisk assigned has insufficient space.

SFW retries the cache area recreation process after an interval of 15 minutes. Within this interval, either decrease your CA size or ensure that the required SSD storage is available on the target ESX.

A cache area recreation is in process in a VMware environment with vMotion, does it affect the sfcache operations?

If the sfcache operations are in process during a vMotion, the sfcache operations fail. You must cancel the operations and retry after the cache area is created.

Does SmartIO continue to use the previous cache area if the VM is moved back to the previous host?

No, moving back the VM to the previous host is considered as a vMotion event. SFW deletes the previous cache area and recreates a new cache area after receiving the first I/O.

This behavior applies even when a virtual machine is powered-on after it was suspended or and it was shutdown for a while.

How does SmartIO behave if the vxsvc service fails during vMotion?

If the vxsvc service fails during vMotion, SFW does not take any action to recreate the cache area. After the vxsvc service is brought online, SFW initiates the cache area recreation process.

Dynamic Multi-Pathing

This chapter includes the following topics:

- [DMP overview](#)
- [Major features of Dynamic Multi-Pathing](#)
- [Active/Active and Active/Passive settings](#)
- [DMP DSMs](#)

DMP overview

The Dynamic Multi-Pathing option adds additional fault tolerance to disk storage by making use of multiple paths between a computer and a disk in a storage array. A path is a connection between the computer and the storage array's disks and consists of a host adapter and a SCSI bus connection to one or more SCSI disks or a fiber optic channel that is connected to a hub, switch, or array. Thus, multiple paths are made possible by connecting two or more host bus adapters with either SCSI or fiber optic cabling to the storage array. The Dynamic Multi-Pathing software manages the multiple paths so that the data on each of the array's disks is accessible to the host computer. If a path to a disk fails, Dynamic Multi-Pathing automatically transmits data to and from the disk over an alternate path.

The paths on an array are set up to work in two ways—either in Active/Active mode, which provides load balancing of the data between multiple paths, or in Active/Passive mode, in which only one path is active and any remaining paths are backups.

SFW offers Dynamic Multi-Pathing as DMP device-specific modules (DSMs).

Note: Windows Management Instrumentation (WMI) is a required service for the MPIO feature to work. Without this service, the MPIO Provider does not load and MPIO objects are not displayed on the VEA console or CLI command outputs. MPIO operations too cannot be performed in absence of WMI service.

DMP DSMs are designed to support a multi-path disk storage environment set up with the Microsoft Multipath I/O (Microsoft MPIO) solution. DMP DSMs work effectively with Windows to provide a fault tolerant multi-path disk storage environment. DMP DSMs have the following benefits:

- Fiber Channel StorPort Miniport HBA Driver support
- iSCSI HBA support
- Microsoft iSCSI Software Initiator support
- Boot from SAN support
- Active/Active Dynamic Multi-Pathing with clustering support
- DMP DSMs load balancing support

Load balancing support includes the following:

- Round robin load balancing support (Active/Active)
- Fail Over Only (Active/Passive)
- Dynamic Least Queue Depth load balancing support
- Round robin with Subset load balancing support
- Least Blocks load balancing support
- Weighted Paths load balancing support
- Balanced Path load balancing support

Basic disks are supported if SCSI-3 registry support is not enabled for DMP DSMs. If SCSI-3 registry is enabled for DMP DSMs for an attached array, then existing basic disks should be upgraded to dynamic disks before placing them under control of DMP DSMs.

For DMP DSMs, Boot and data volumes are supported on the same bus/HBAs for non-clustered servers if the Boot from SAN recommendations from Microsoft is followed. DMP DSMs are not supported with Fibre Channel port drivers, Fibre Channel SCSI Miniport drivers or boot or cluster disks on the same bus/HBAs.

Major features of Dynamic Multi-Pathing

The major features of Dynamic Multi-Pathing (DMP) are described in this section.

They include the following:

- **Fault tolerance**
Provides fault tolerance to a disk system by using multiple paths to each disk. If the primary path fails, either at the card level or in the cabling from the card to the disk, a secondary path is automatically used.
- **Load balancing in Active/Active configurations**
When a system is configured as Active/Active, Dynamic Multi-Pathing makes use of all the paths to a disk for the transfer of I/O to and from the disk.
- **Support for multiple paths**
With DMP DSMs, the maximum number of I/O paths you can have is 16 per array LUN.
- **Dynamic recovery**
If an active path to a disk fails, Dynamic Multi-Pathing automatically flags the failed path and no longer attempts to transfer data on it. The failed path is monitored and is automatically restored to service when Dynamic Multi-Pathing detects that the path functions correctly. Dynamic Multi-Pathing automatically updates path status on the user display when a path fails or is restored to service.
- **Dynamic path recognition**
If you add a new path to your Dynamic Multi-Pathing configuration, running a rescan or restarting your system causes Dynamic Multi-Pathing to detect the new path and display its status. If a failed or disabled path is restored to service, Dynamic Multi-Pathing automatically detects the status change and updates the display.

Active/Active and Active/Passive settings

Dynamic Multi-Pathing has two modes of operation for an array's paths, Active/Active and Active/Passive.

These modes also apply to the array's disks and are defined as follows:

- **Active/Active**
The mode in which Dynamic Multi-Pathing allocates the data transfer across all the possible paths, thus enabling the desirable feature of load balancing. With this mode, Dynamic Multi-Pathing implements a round-robin algorithm, selecting each path in sequence for each successive data transfer to or from a disk. For example, if you have two paths active, A and B, the first disk transfer occurs on path A, the next on path B, and the next on path A again.

In addition to the round-robin algorithm, DMP DSMs offer the following load balancing options:

- Dynamic Least Queue Depth

Selects the path with the least number of I/O requests in its queue for the next data transfer.

For example, if you have two active paths, path A with one I/O request and path B with none, DMP DSMs selects the path with the least number of I/O requests in its queue, path B, for the next data transfer.

- Balanced Path

This policy is designed to optimize the use of caching in disk drives and RAID controllers. The size of the cache depends on the characteristics of the particular hardware. Generally, disks and LUNs are logically divided into a number of regions or partitions. I/O to and from a given region is sent on only one of the active paths. Adjusting the region size to be compatible with the size of the cache is beneficial so that all the contiguous blocks of I/O to that region use the same active path. The value of the partition size can be changed by adjusting the value of the tunable parameter, Block Shift.

Block Shift represents the number of contiguous I/O blocks that are sent along a path to an Active/Active array before switching to the next available path. The Block Shift value is expressed as the integer exponent of a power of 2. For example, the Block Shift value of 11 represents 2^{11} or 2048 contiguous blocks of I/O.

The benefit of this policy is lost if the value is set larger than the cache size. The benefit is also lost when the active path fails. In this situation, the I/O is automatically redistributed across the remaining paths.

The default value of the Block Shift parameter is set to 11 so that 2048 blocks (1MB) of contiguous I/O are sent over a path before switching to a different path. Depending on your hardware, adjusting this parameter may result in better I/O throughput. Refer to your hardware documentation for more information.

Note: Block Shift only affects the behavior of the balanced path policy. A value of 0 disables multi-pathing for the policy unless the `vxsdmpadm` command is used to specify a different partition size for an array.

- Weighted Paths

Uses the path with the lowest numerical weight. Each path is assigned a weight by the user to designate which path is favored for data transfer. If two or more paths have the same weight and are the lowest weight of all paths, then these paths are used each in turn, in round-robin fashion, for the data transfer.

For example, if you have three active paths, path A with weight of 0, path B with weight of 0, and path C with weight of 9, DMP DSMs use path A for one

data transfer and then use path B for the next. Path C is in standby mode and is used if path A or path B fails.

- Round robin with Subset
Uses a subset of paths, each in turn, in round-robin fashion. The user specifies the paths for data transfer that make up the subset. The remaining paths are in standby mode.
For example, if you have three active paths, path A, path B, and path C and you specify the subset to contain path A and path B, then DMP DSMs use path A for one data transfer and then use path B for the next. Path C is in standby mode and is used if path A or path B fails.
- Least Blocks
Selects the path with the least number of blocks of I/O in its queue for the next data transfer.
For example, if you have two active paths, path A with one block of I/O and path B with none, DMP DSMs select the path with the least number of blocks of I/O in its queue, path B, for the next data transfer.
- Active/Passive
A mode in which a path that is designated as the "Preferred Path" or "primary path" is always active and the other path or paths act as backups (standby paths) that are called into service if the current operating path fails.

The modes of operation—Active/Active and Active/Passive—are shown as options in the Load Balancing section of the program's Array Settings and Device Settings windows. The Active/Active mode enables load balancing, but the Active/Passive mode does not provide load balancing except for the Fail Over Only load balancing policy.

Note: If a storage array cannot transfer data on one of the path configurations, the Load Balancing options appear grayed out on the screen and you cannot access these settings.

You configure the load balancing settings for the paths at the array level through the Array Settings screen, or you can accept the default setting. The default setting is dependent on the particular array. Consult the documentation for your storage array to determine the default setting of the array and any additional settings it supports.

After the appropriate array setting is made, all the disks in an array have the same load balancing setting as the array. If the array is set to active/active, you can use the Device Settings screen to change the setting on an individual disk so that it has a different load balancing setting than the array. When an array is set to

active/passive, no load balancing is enabled and data transfer is limited to the one preferred or primary path only.

For all Active/Active arrays under control of DMP DSMs:

- All paths to the disks are current active I/O paths. Each active path is designated by a path icon with a green circle in the VEA GUI.
- For an Active/Passive load balance setting, the primary path is designated by a path icon with a checkmark in a green circle in the GUI.
- The DMP DSMs are not enabled to indicate which array controller each path is connected to.

For all Active/Passive Concurrent (A/PC) and Asymmetric Logical Unit Access (ALUA) arrays under control of DMP DSMs, the load balance settings apply only to the current active I/O paths. If all the active I/O paths change or fail, the load balance settings are automatically applied to the new current active I/O paths of the arrays.

In addition, for A/PC and ALUA arrays:

- The current active path is designated by a path icon with a green circle in the VEA GUI.
- For an Active/Passive load balance setting, the primary path is designated by a path icon with a checkmark in a green circle in the VEA GUI.
- DMP automatically selects the primary path for Active/Passive load balancing.
- Round robin with Subset and Weighted Paths load balance settings are available only at the device level. They are not available at the array level.
- Active paths are connected to the same array controller.

Active/Active and Active/Passive settings in a cluster environment

This section covers information about settings for DMP DSMs along with information about enabling or disabling SCSI-3 PGR.

DMP DSMs

For DMP DSMs in a cluster environment, either Active/Active or Active/Passive load balance settings can be used. DMP DSMs automatically set the load balancing to active/passive for disks under SCSI-2 reservation. For Active/Active load balancing in a cluster environment, the array must be enabled for SCSI-3 Persistent Group Reservations (SCSI-3 PGR).

Note: Veritas maintains a Hardware Compatibility List (HCL) for Storage Foundation and High Availability Solutions for Windows products on the Veritas support Web site. The HCL gives information on HBAs, firmware, and switches that have been tested with each supported array. Check the HCL for details about your hardware before using DMP DSMs.

Storage arrays may require additional configuration steps or updates to work with Storage Foundation for Windows and MPIO. Contact the manufacturer of the storage array for details.

SCSI-3 PGR technology

SCSI-3 PGR supports multiple nodes accessing a device while at the same time blocking access to other nodes. SCSI-3 PGR supports multiple paths from a host to a disk and SCSI-3 PGR reservations are persistent across SCSI bus resets.

By contrast, SCSI-2 reservations support only one host, with one path. This means if there is a need to block access for data integrity concerns, only one host, and one path remain active. This limits access to the device to only one path and prevents the use of multiple paths even if they are available.

SCSI-3 PGR uses a concept of registration and reservation. Systems accessing a device register a key with a SCSI-3 device. Each system registers its own key. Multiple systems registering keys form a membership. Registered systems can then establish a reservation. The reservation type is set to "Exclusive Access - Registrants Only". This means that only some commands are allowed and there is only one persistent reservation holder. With SCSI-3 PGR technology, blocking write access can be done by removing a registration from a device.

In the Storage Foundation implementation, a node registers the same key for all paths to the device.

For DMP DSMs, the Active/Active setting is implemented by translating SCSI reserve or release commands to SCSI-3 PGR commands.

Enabling SCSI-3 PGR

You must enable SCSI-3 support before using the Active/Active setting for DMP DSMs. (SCSI-3 support is disabled by default.)

Note: Before enabling SCSI-3 PGR support, move any cluster disk groups to another node or deport the cluster disk groups.

In SFW 5.1 and later releases, SFW support for clustering environments can be selected to issue SCSI-2 or SCSI-3 reservations with the SFW Control Panel.

- If SCSI-3 is selected in the SFW Control Panel, then there is no need to enable SCSI-3 PGR support for the DMP DSMs.

After selecting SCSI-3, you must issue the following CLI commands to complete the setting:

- `net stop vxsvc`
- `net start vxsvc`

- If SCSI-2 is selected in the SFW Control Panel, then SCSI-3 PGR support can be enabled for the DMP DSMs by updating the Windows registry by running one of the files that is listed in the table below, or by using the `vxdkmpadm` CLI command.

See “[vxdkmpadm](#)” on page 690.

See “[Additional considerations for SFW Microsoft Failover Clustering support](#)” on page 503.

The files in the following table are provided to enable or disable SCSI-3 support for the type of storage array supported. The files are located on the product CD at `..\Tools\storage_foundation_for_windows\DMP_DSM_SCSI3_reg`

Copy the files to your system and run the appropriate `.reg` for your storage array.

Note: You must ensure that your storage array supports SCSI-3 PGR before using the Active/Active setting for DMP DSMs. Also check to see if your array manufacturer requires any special settings for a LUN to be enabled for SCSI-3 PGR.

Table 13-1 Files to Enable/Disable SCSI-3 PGR

Array Type	Enable SCSI-3 PGR	Disable SCSI-3 PGR
<ul style="list-style-type: none"> ■ EMC Symmetrix 8000 Series ■ EMC Symmetrix DMX Series 	EnableSupportSCSI3EMC.reg	DisableSupportSCSI3EMC.reg
<ul style="list-style-type: none"> ■ IBM TotalStorage™ ESS800/ESS750 ■ IBM TotalStorage™ DS8000 	EnableSupportSCSI3IBMDS.reg	DisableSupportSCSI3IBMDS.reg

Table 13-1 Files to Enable/Disable SCSI-3 PGR (*continued*)

Array Type	Enable SCSI-3 PGR	Disable SCSI-3 PGR
<ul style="list-style-type: none"> ■ Hitachi TagmaStore Universal Storage Platform (USP100, USP600, and USP1100) ■ Hitachi TagmaStore Network Storage Controller (NSC55) ■ Hitachi TagmaStore Adaptable Modular Storage (AMS200 and AMS500) ■ Hitachi TagmaStore Workgroup Modular Storage (WMS100) ■ Hitachi Hitachi 9900 Lightning Series (9900 and 9900V) ■ Sun StorEdge SE9900 Series (SE9910, SE9960, SE9970V, SE9980V, and SE9990) 	EnableSupportSCSI3HDS.reg	DisableSupportSCSI3HDS.reg
<ul style="list-style-type: none"> ■ HP StorageWorks XP128 Disk Array ■ HP StorageWorks XP1024 Disk Array ■ HP StorageWorks XP10000/XP12000 Disk Array 	EnableSupportSCSI3HPXP.reg	DisableSupportSCSI3HPXP.reg
<ul style="list-style-type: none"> ■ HP StorageWorks Enterprise Virtual Array (EVA4000, EVA6000, and EVA8000) 	EnableSupportSCSI3HPEVA.reg	DisableSupportSCSI3HPEVA.reg
<ul style="list-style-type: none"> ■ Network Appliance F800 Series ■ Network Appliance FAS200 Series (FAS250 and FAS270) ■ Network Appliance FAS900 Series (FAS920, FAS940, FAS960, and FAS980) ■ Network Appliance FAS3000 Series (FAS3020 and FAS3050) ■ Network Appliance NearStore Series ■ NetApp V-Series (GF980c, GF960c, V3050c, V3020c, and GF270c) <p>Note: The only Network Appliance arrays that are supported are those that are configured as an active/active hardware type (cfmode: standby for each Network Appliance filer).</p>	EnableSupportSCSI3NETAPP.reg	DisableSupportSCSI3NETAPP.reg

DMP DSMs

This section describes support for the Microsoft MPIO solution. DMP DSMs provide the support.

This section covers the following topics:

- See [“DMP DSMs menus”](#) on page 408.
- See [“Add and remove paths”](#) on page 412.
- See [“Specify load balancing settings and the primary path”](#) on page 413.
- See [“View array, disk, and path status”](#) on page 421.
- See [“DMP DSMs”](#) on page 407.

DMP DSMs menus

This section describes DMP DSMs menus.

They include the following:

- See [“Displaying the menu for arrays”](#) on page 408.
- See [“Commands for arrays”](#) on page 408.
- See [“Displaying the menu for disks”](#) on page 409.
- See [“Commands for paths of a disk”](#) on page 410.

Displaying the menu for arrays

There is a context menu for arrays that are under the control of DMP DSMs that let you access the settings and properties of the array. It is available when you select an array in the VEA GUI.

To display the menu for arrays

- 1 Open the DMP DSMs folder in the VEA GUI.
- 2 Select the array that you are interested in and the array context menu appears.

Commands for arrays

This section describes the DMP DSMs menu commands.

Array settings

The Array Settings command launches the Array Settings wizard.

It lets you select the setting for the operational modes for the array's paths in the Load Balance Policy section. The Control Timer Settings is a section of tunable parameters for testing the status or health of a path.

At the bottom of the screen there is an option for applying all the settings to all the arrays under control of DMP DSMs. When you select this option, the current settings for this array are applied to all the arrays that the specific DMP DSMs driver controls.

Note: There are specific DMP DSMs that work with specific families of hardware arrays. If your environment has more than one family of hardware arrays under control of DMP DSMs, then each family of hardware arrays work with specific DMP DSMs. When you select the option for applying all the settings to all the arrays under control of the DSM, you apply the settings to only those arrays that work with that specific DMP DSMs.

See the Veritas Web site for more information about which DMP DSMs support which hardware arrays.

Additional information about array settings is available.

See [“Array status”](#) on page 421.

Note: If the Fail Over Only (Active/Passive) array load balancing setting is selected and a primary path has not been set, DMP DSMs set the primary path. Generally, DMP DSMs set the primary path to the first path that is able to transfer data. You can manually set the primary path with the Set primary path command.

Properties

The Properties command brings up the Properties screen, which displays information about the selected array, including array name, array type, load balance policy, and a list of the devices in the array.

Additional information about the array, disk, and path settings is available.

See [“View array, disk, and path status”](#) on page 421.

Displaying the menu for disks

The context menu for a disk is available when you highlight a path that is connected to the disk. The menu requires the DMP DSMs tab to be active.

To display the disk context menu

- 1 Open the DMP DSMs folder in the tree view.
- 2 Select the disk that you are interested in.

Note that opening the Disks folder and clicking a disk that is under DMP DSMs control is an alternative way to select a disk.

- 3 Select the DMP DSMs tab in the right pane. This tab displays the paths to the selected disk.

Each path is displayed with its properties. Some of the attributes that are displayed are Port, Channel, Target ID, LUN, and Load Balance Policy.

The name of the path is derived from a combination of properties of the path. For example, if the path has properties of port number = 2, Channel = 1, and Target ID = 1, then the name of the path is 2-1-1.

- 4 Right-click on a path and the context menu appears.

Commands for paths of a disk

This section describes the DMP DSMs menu commands.

The menu commands are the following:

See [“Statistics monitoring”](#) on page 410.

See [“Clear performance statistics”](#) on page 410.

See [“Device settings”](#) on page 411.

See [“Set primary path”](#) on page 411.

See [“Properties”](#) on page 412.

Statistics monitoring

The Statistics Monitoring command brings up the Path Statistics screen, which displays the I/O statistics of the selected path.

The I/O statistics for the selected path are the following:

- Number of read operations per second
- Number of write operations per second
- Number of bytes read per second
- Number of bytes written per second

The statistics are updated and displayed at a specified polling interval. The range of values for the polling interval is 2 to 900 seconds. You can set the polling interval by entering the number of seconds in the Polling Interval field and clicking **Set Polling Interval**.

Clear performance statistics

The Clear Performance Statistics command brings up the Clear Performance Statistics screen, which lets you clear the performance statistics that are monitored

for a single device, an array, or for all the arrays under control of a DMP DSM. After selecting which performance statistics to clear, click **OK** to complete the operation.

Device settings

The Device Settings command brings up the Device Settings screen, which displays the DMP DSMs characteristics of the disk whose path is selected on the DMP DSMs tab. It has a Load Balance Policy section to set the operational mode that applies to the selected disk.

See [“Device status”](#) on page 422.

Note: If the Fail Over Only (Active/Passive) device load balancing setting is selected and a primary path has not been set, DMP DSMs set the primary path. Generally, DMP DSMs set the primary path to the first path that is able to transfer data. You can manually set the primary path with the Set primary path command.

Set primary path

The Set primary path command on the DMP DSMs menu can be used to specify the currently selected path as the primary path to the disk or the array.

The following options are available:

- Set primary path for device
This menu option is available only when the operational mode for the path to the disk is specified as Fail Over Only (Active/Passive) and the currently selected path is not the primary path. The primary path is the only path that is used to transfer data to and from the disk. Clicking this menu selection makes the selected path the primary path to the disk.
See [“Specifying the primary path for an array or disk”](#) on page 418.
- Set primary path for array
This lets you specify the currently selected path as the primary path to the array. The primary path is the only path that is used to transfer data to and from the array.
The array's primary path can be set regardless of the array's current load balance setting. This means that if a disk in the array is set to Fail Over Only (Active/Passive), the primary path for the array becomes the primary path for the disk. If the setting is set to any other load balance option, then the primary path setting for the array is not used. However, if a disk in the array is set to Fail Over Only (Active/Passive), the primary path for the array becomes the primary path for the disk.
For example, if an array has a load balance setting of round robin (Active/Active) and the primary path for the array is set to path 2-1-1, then any disk in the array

that has the load balance setting of Fail Over Only (Active/Passive) has its primary path set to path 2-1-1.

See [“Specifying the primary path for an array or disk”](#) on page 418.

Properties

The Properties command brings up the Properties screen, which displays information about the selected path, including path name, status, and SCSI address.

See [“Path status”](#) on page 422.

Add and remove paths

After using DMP DSMs for a while, you might want to add a path to an array if you want to add more disks to the array or if you want additional fault tolerance for the array by having another path. In addition, you may want to remove a path if you are reconfiguring your system.

This section describes how to add and remove paths.

See [“Adding a path”](#) on page 412.

See [“Removing a path”](#) on page 413.

Adding a path

You might want to add a path when you add more disks to an array or to increase fault tolerance for the array.

To add a path to an array

- 1 Add the path by physically connecting a new cable.
- 2 To view the screen changes caused by the additional path, open the Disks folder in the tree view and select a disk. Then select the DMP DSMs tab in the right-hand pane.
- 3 Select **Actions > Rescan** from VEA.

DMP DSMs scans for paths and updates the screen, reflecting any changes you have made in the physical installation.

Note: If a disk shares a path with another disk, then it must share all the same paths with the other disk. If a disk in the array does not share a path with any other disk in the array, then all the paths of the disk must not be shared with any other disk in the array.

Removing a path

Remove a path by physically disconnecting the cable. The system monitors existing paths and detects that the disconnected path is no longer functioning properly. The SFW display is updated to indicate this, and data transfers failover to the next path in the sequence.

Note: You may need to perform a rescan to update the screen after removing a path from an array that was under control of DMP DSMs. Select **Actions > Rescan** from VEA to display the change you have made to the configuration.

Specify load balancing settings and the primary path

The primary path setting is only available when you select the Fail Over Only (Active/Passive) load balancing setting for an array or disk.

When you first set up an array under DMP DSMs, you must make sure that you have the load balancing setting you want for the paths in the array. After that is done, all of the disks in the array by default have the same load balancing setting. Then, if desired, you can change this setting for individual disks. The load balancing settings are located in the Load Balance Policy section of the Array Settings wizard and the Device Settings screen.

All of the load balancing settings, except for Fail Over Only (Active/Passive), enable load balancing among the active paths. The Fail Over Only (Active/Passive) setting does not provide load balancing because data transfer only occurs on the one active path, the primary path.

Note: The type of array or disk determines which load balancing settings are available. Load balance settings that are not supported are grayed-out in the VEA GUI. Refer to the documentation for your array or disk for more information about supported load balance settings.

This section gives the step-by-step directions for specifying the load balancing settings for the paths in an array and for individual disks.

See [“Specifying load balancing settings for an array”](#) on page 414.

See [“Specifying load balancing settings for a disk”](#) on page 416.

See [“Specifying the primary path for an array or disk”](#) on page 418.

See [“Active/Active and Active/Passive settings”](#) on page 401.

Specifying load balancing settings for an array

The way to specify load balancing settings using the Array Settings wizard follows.

To specify load balancing settings for an array

- 1 To launch the Array Settings wizard, open the DMP DSMs folder in the tree view.
- 2 Right-click the array that you are interested in and the array context menu appears.
- 3 Select **Array Settings** from the context menu.
- 4 The Array Settings welcome screen appears. Click **Next** to continue.
- 5 The Select Array Settings parameters screen appears.

Select the load balancing option by clicking the appropriate radio button.

SCSI-3 reservation support is required for all load balancing settings, except for Fail Over Only (Active/Passive). Refer to your array documentation to determine if your array supports SCSI-3 reservations before selecting any of these load balancing settings.

- Round robin (Active/Active)
This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths in the array as active and I/O transfers occur on the paths each in turn, in a round-robin fashion.
- Dynamic Least Queue Depth
This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the least number of I/O requests in its queue for a data transfer.
- Balanced Path
This policy is designed to optimize the use of caching in disk drives and RAID controllers. The size of the cache depends on the characteristics of the particular hardware. Generally, disks and LUNs are logically divided into a number of regions or partitions. I/O to and from a given region is sent on only one of the active paths. Adjusting the region size to be compatible with the size of the cache is beneficial so that all the contiguous blocks of I/O to that region use the same active path. The value of the partition size can be changed by adjusting the value of the tunable parameter, Block Shift.

More details about this option are available.

See [“Active/Active and Active/Passive settings”](#) on page 401.

- **Weighted Paths**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the lowest weight for data transfer. The user assigns weights to designate which path is favored. If two or more paths have the same weight and are the lowest weight of all paths, then these paths are used each in turn, in round-robin fashion, for the data transfer.

Note: This setting is available for A/PC and ALUA arrays only at the device level. It is not available at the array level.

- **Round robin with Subset**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and uses a subset of paths. These paths are used each in turn, in round-robin fashion, for data transfer. The user specifies the subset of paths. The remaining paths are in standby mode.

Note: This setting is available for A/PC and ALUA arrays only at the device level. It is not available at the array level.

- **Least Blocks**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the least number of blocks of I/O in its queue for data transfer.

- **Fail Over Only (Active/Passive)**

This option has one primary path with the remaining path or paths on standby (backups) in case the primary path fails. The primary path is the only path that is used to transfer data. This option does not provide load balancing among paths. DMP DSMs sets a path to be the primary path. However, you may choose a specific path to be the primary path.

See [“Specifying the primary path for an array or disk”](#) on page 418.

There are specific DMP DSMs that work with specific families of hardware arrays. If your environment has more than one family of hardware arrays under control of DMP DSMs, then each family of hardware arrays works with specific DMP DSMs.

If you want the selected load balancing setting to be the setting for all the arrays that this specific DMP DSMs driver controls, check the check box for **All arrays under the DSM**.

The default load balancing setting of an array is dependent on the particular array. Consult the documentation for your storage array to determine the default setting and any additional settings it supports.

In the sample screen that is shown below, the round robin (Active/Active) setting is selected.

Click **Next** to continue.

- 6 The Array Settings summary screen appears displaying the settings for the array. Review the settings and click **Finish** to apply the settings and exit the wizard.

For some load balance policies, such as round robin with Subset, additional specifications are required. The Array Settings wizard displays a screen before the summary screen to collect these specifications. Click **Next** after setting these specifications to continue to the summary screen.

Specifying load balancing settings for a disk

If a setting has not been specified for an individual disk, the disk assumes the same load balancing setting as the one used by the array. However, by using the Device Settings screen, you can change the load balancing option for an individual disk and make it different than the setting of the array.

DMP DSMs automatically set the load balancing to Fail Over Only (Active/Passive) for disks under SCSI-2 reservation. The Fail Over Only (Active/Passive) setting is always used on disks under SCSI-2 reservation and the system ignores other load balance settings that may be set.

Disks under SCSI-2 reservation are:

Disks that may be in a clustering environment.

Disks with Private Disk Group Protection.

To specify load balancing settings for a disk

- 1 Open the DMP DSMs folder in the VEA console's tree view.
- 2 Select the disk that has the load balancing settings you want to change.

Note that opening the Disks folder and clicking a disk that is under DMP DSMs control is an alternate way to select a disk.

- 3 Make sure the DMP DSMs tab is selected in the right pane, right-click one of the displayed paths, and then select **Device Settings** from the path context menu.
- 4 The Device Settings welcome screen appears. Click **Next** to continue.
- 5 The Select Device Settings parameters screen appears.

Select one of the following load balancing options:

- **Round robin (Active/Active)**

This option ensures that load balancing is used for data transfers to and from a disk. With this setting, DMP DSMs configures all paths to the disk as active and enables paths, each in turn, in a round-robin fashion for data transfer.
- **Dynamic Least Queue Depth**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the least number of I/O requests in its queue for a data transfer.
- **Balanced Path**

This policy is designed to optimize the use of caching in disk drives and RAID controllers. The size of the cache depends on the characteristics of the particular hardware. Generally, disks and LUNs are logically divided into a number of regions or partitions. I/O to and from a given region is sent on only one of the active paths. Adjusting the region size to be compatible with the size of the cache is beneficial so that all the contiguous blocks of I/O to that region use the same active path. The value of the partition size can be changed by adjusting the value of the tunable parameter, Block Shift.

More details about this option are available.
See [“Active/Active and Active/Passive settings”](#) on page 401.
- **Weighted Paths**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the lowest weight for data transfer. The user assigns weights to designate which path is favored. If two or more paths have the same weight and are the lowest weight of all paths, then these paths are used each in turn, in round-robin fashion, for the data transfer.
- **Round robin with Subset**

This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array

as active and uses a subset of paths each in turn, in round-robin fashion, for data transfer. The user specifies the subset of paths. The remaining paths are in standby mode.

- **Least Blocks**
This option ensures that load balancing is used for data transfers to and from an array. With this setting, DMP DSMs configures all paths to the array as active and selects the path with the least number of blocks of I/O in its queue for data transfer.
- **Fail Over Only (Active/Passive)**
This option has one primary path with the remaining path or paths on standby (backups) that are used when the current primary path fails. The primary path is the only path that is used to transfer data. This option does not provide load balancing among paths. DMP DSMs chooses a path to be the primary path. However, you may choose a specific path to be the primary path.
See [“Specifying the primary path for an array or disk”](#) on page 418.

Click **Next** to continue.

- 6 The Device Settings summary screen appears displaying the settings for the array. Review the settings and click **Finish** to apply the settings and exit the wizard.

For some load balance policies, such as round robin with Subset, additional specifications are required. The Array Settings wizard displays a screen before the summary screen to collect these specifications. Click **Next** after setting these specifications to continue to the summary screen.

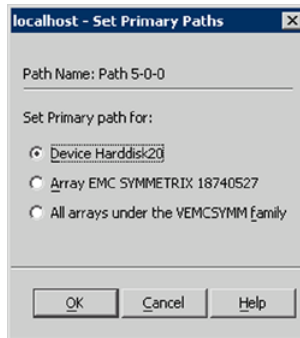
Specifying the primary path for an array or disk

When an array is configured for Fail Over Only (Active/Passive) operation, you can specify the path that is used for data transfer. The specified path, called the primary path, is the only path used for data transfer.

To specify the primary path for an array or disk

- 1 Display the Set primary path screen by doing the following:
 - Make sure the DMP DSMs tab is selected in the right pane.
To display the DMP DSMs tab, use the tree view in the left pane to select a disk under the DMP DSMs folder, or a disk under the Disks folder that has DMP DSMs enabled, and then click the DMP DSMs tab that appears in the right pane.
 - In the DMP DSMs tab, right-click on the path you want to be the primary path for the array, and select **Set Primary Path** from the path context menu.

The Set primary path screen appears.



2 Select to set the primary path for the device or for the array, as follows:

■ Primary path for device

The primary path for the device is available only when the operational mode for the path to the disk is specified as Fail Over Only (Active/Passive) and the currently selected path is not the primary path. The primary path is the only path that is used to transfer data to and from the disk. Clicking this menu selection makes the selected path the primary path to the disk. The other paths to the disk are placed on standby and available for failover.

■ Primary path for array

This lets you specify the currently selected path as the primary path to the array. The primary path is the only path that is used to transfer data to and from the array.

The array's primary path can be set regardless of the array's current load balance setting. This means that if the load balance setting for the array is set to round robin (Active/Active) or other Active/Active load balance setting, then the primary path setting for the array is not used. However, if a disk in the array is set to Fail Over Only (Active/Passive), the primary path for the array becomes the primary path for the disk.

For example, if an array has a load balance setting of round robin (Active/Active) and the primary path for the array is set to path 2-1-1, then any disk in the array that has the load balance setting of Fail Over Only (Active/Passive) has its primary path set to path 2-1-1.

3 Click **OK** to save your setting and exit the screen.

A blue checkmark icon identifies the primary path.

Specifying control timer settings for an array

Control Timer Settings for an array are a set of tunable parameters that affect the testing of a path's status or health.

To specify control timer settings for an array

- 1 To launch the Array Settings wizard, open the DMP DSMs folder in the tree view.
- 2 Right-click the array that you are interested in and the array context menu appears.
- 3 Select **Array Settings** from the path context menu.
- 4 The Array Settings welcome screen appears. Click **Next** to continue.
- 5 The Select Array Settings parameters screen appears.

Edit the values for the parameters.

Click **Next** to continue.

The Control Timer Settings parameters are as follows:

Test Path Retry Count	The number of times DMP DSMs test a path to determine if it has recovered from an error condition. The default value is 0. The maximum value is 50.
SCSI Command Timeout	The amount of time a SCSI command waits for a device to respond to it. The default value is 30 seconds. The maximum value is 120 seconds.
Kernel Sleep Time	The interval of time the DMP DSMs kernel waits between attempts to recover a path from an error condition. The default value is 200 milliseconds. The maximum value is 10000 milliseconds.
Failover Retry Count	The number of times DMP DSMs attempt a failover to another path when the current path fails. The default value is 0. The maximum value is 50.

- 6 The Array Settings summary screen appears displaying the settings for the array. Review the settings and click **Finish** to apply the settings and exit the wizard.

For some load balance policies, such as round robin with Subset, additional specifications are required. The Array Settings wizard displays a screen before the summary screen to collect these specifications. Click **Next** after setting these specifications to continue to the summary screen.

View array, disk, and path status

DMP DSMs provide status information for arrays, disks, and paths to the disks. Status displays are accessed through the path context menu.

See [“Array status”](#) on page 421.

See [“Device status”](#) on page 422.

See [“Path status”](#) on page 422.

Array status

To view array status with Array Properties, perform the following steps:

To view array status

- 1 Open the DMP DSMs folder in the tree view.
- 2 Right-click the array that you are interested in and the array context menu appears.
- 3 Select **Properties** from the context menu.
- 4 The Properties screen appears.

The Properties screen contains the following:

Array Name	Displays the name of the array. The name is not user definable.
Array Type	Displays the type of load balance policy in effect for the array.
Control Timer Settings	<p>A set of the following tunable parameters that affect the testing of a path’s status or health.</p> <ul style="list-style-type: none">■ Test Path Retry Count — The number of times DMP DSMs test a path to determine if it has recovered from an error condition. The default value is 0. The maximum value is 50.■ SCSI Command Timeout — The amount of time a SCSI command waits for a device to respond to it. The default value is 30 seconds. The maximum value is 120 seconds.■ Kernel Sleep Time — The interval of time the DMP DSMs kernel waits between attempts to recover a path from an error condition. The default value is 200 milliseconds. The maximum value is 10 seconds.■ Failover Retry Count — The number of times DMP DSMs attempt a failover to another path when the current path fails. The default value is 0. The maximum value is 50. Click OK to exit.

Load Balance Policy	<p>Displays the load balancing setting in effect. There is one option, Fail Over Only (Active/Passive), that does not provide load balancing. All other options provide load balancing.</p> <p>Option descriptions are available.</p> <p>See “Active/Active and Active/Passive settings” on page 401.</p> <p>The procedure to specify settings is described also.</p> <p>See “Specifying load balancing settings for an array” on page 414.</p>
Devices in the Array	<p>Displays the number of disks in the array and lists them by name.</p>

Device status

The way to view device status is as follows.

To view device status

- 1 In the tree view under the DMP DSMs folder, select a disk whose status you want to view.
- 2 In the right pane, click the DMP DSMs tab for the disk.
- 3 Right-click the path and select **Device Settings** from the path context menu that comes up.
- 4 The Device Settings welcome screen appears. Click **Next** to continue.

The Device Settings screen appears.

The Device Settings screen contains the setting for the **Load Balance Policy** for the disk.

Load Balance Policy — Indicates the setting in effect for the individual disk. The Fail Over Only (Active/Passive) option provides no load balancing. The other options provide load balancing.

Option descriptions are available.

See [“Active/Active and Active/Passive settings”](#) on page 401.

The procedure to specify settings is also described.

See [“Specifying load balancing settings for a disk”](#) on page 416.

- 5 Click **Cancel** to exit.

Path status

The way to view path status is as follows.

To view path status

- 1 Make sure that the path you want to view the status for is displayed in the DMP DSMs tab in the right pane of the VEA console.

If the DMP DSMs tab is not displayed, select a disk that uses that path in the tree view under the Disks folder. Then, in the right pane, click the DMP DSMs tab for the disk.

- 2 Select the path in the DMP DSMs tab.

Select **Properties** from the path context menu that comes up.

The Path Properties screen appears.

The screen displays the following information for the selected path:

Name	The name of the selected path.
Status	DMP DSMs have two path state indicators: <ul style="list-style-type: none">■ Healthy — The path is operational.■ Unhealthy — The path is not operational.
Primary	Whether the path is primary or not.
Port Number	The number of the physical port on the storage array that the selected path is connected to.
Channel	The channel that is associated with the path.
Target ID	The number that, when combined with the LUN, uniquely identifies a disk on the port.
LUN	Logical Unit Number, the number that, when combined with the Target ID, uniquely identifies a disk on the port.
Load Balancing Policy	Current load balance policy in effect.

- 3 Click **OK** to close the screen.

Configuring Cluster Volume Manager (CVM)

This chapter includes the following topics:

- [Overview](#)
- [Configuring a CVM cluster](#)
- [Administering CVM](#)
- [Unconfiguring a CVM cluster](#)
- [Command shipping](#)
- [About I/O Fencing](#)

Overview

Cluster Volume Manager (CVM) is a new way to do storage management in a clustered environment. With CVM, failover capabilities are now available at a volume-level granularity. Volumes under CVM allow exclusive write access across multiple nodes of a cluster. In a Microsoft Failover Clustering environment, you can now create clustered storage out of shared disks, which lets you share volume configurations and enable fast failover support at volume-level. Each node recognizes the same logical volume layout and, more importantly, the same state of all volume resources. The same logical view of disk configuration and any changes to this view are available on all the nodes.

Note: CVM (and related cluster-shared disk groups) is supported only in a Microsoft Hyper-V environment. It is not supported for a physical environment.

The CVM Master node exists on a per-cluster basis and uses GAB (Group Atomic Broadcast) and LLT (Low Latency Transport) to transport its configuration data.

CVM is based on a “Master and Slave” architecture pattern. At any given time, one node of the cluster acts as a Master, while the rest of the nodes take the role of a Slave. The Master node is responsible for maintaining the volume management configuration information. Each time a Master node fails, a new Master node is selected from the surviving nodes.

Configuring a CVM cluster

Configuring a Cluster Volume Manager (CVM) cluster involves the following two tasks:

- Preparing a cluster for CVM using either the GUI-based SFW Configuration Wizard for Microsoft Failover Cluster or the CLI-based `ffconfigcli` utility.

Note: The wizard or the utility servers dual purpose. It is used for configuring CVM for a Hyper-V environment. It is also used to enable fast failover in Microsoft Failover Clustering for a physical environment.

If the nodes participating in the CVM cluster are on multi-subnet cluster networks, then you need to use the CLI-based `SFWConnectionManager` utility to configure a CVM link for the nodes. For more information, See “[Configuring CVM links for multi-subnet cluster networks](#)” on page 433.

- Creating a cluster-shared disk group (CSDG) and cluster-shared volume.

The SFW Configuration Wizard for Microsoft Failover Cluster sets up a configuration for CVM cluster to communicate over the underlying Microsoft Failover Clustering networks.

A dynamic volume that is created under a CSDG is automatically configured in the Microsoft failover cluster as a Volume Manager Shared Volume resource.

Note: For CVM clusters, Veritas recommends that you use a physical disk as the quorum resource.

Note: Before configuring CVM, ensure that the Volume Manager Shared Volume resource has been registered. If not, then you need to register it manually. For information about registering the resource manually, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

Preparing SFW for Microsoft Failover Clustering

This section provides information about preparing a cluster for CVM using the GUI-based SFW Configuration Wizard for Microsoft Failover Cluster and the CLI-based `ffconfigcli` utility. This prepares SFW for a Microsoft Failover Clustering environment.

Please note the following important information regarding using the wizard or utility:

- All the nodes that participate in the CVM cluster must already be part of a Microsoft failover cluster.
- InfoScale Storage must be installed on all the nodes participating in CVM.

Note: InfoScale Enterprise is not supported for configuring CVM.

- The Microsoft failover cluster service (`ClusSvc`) should be running on all the nodes in the cluster.
- If you make any changes to the cluster network configuration of the Microsoft failover cluster after configuring CVM, then you need to unconfigure the CVM cluster and reconfigure it using either the wizard or the `ffconfigcli` utility.

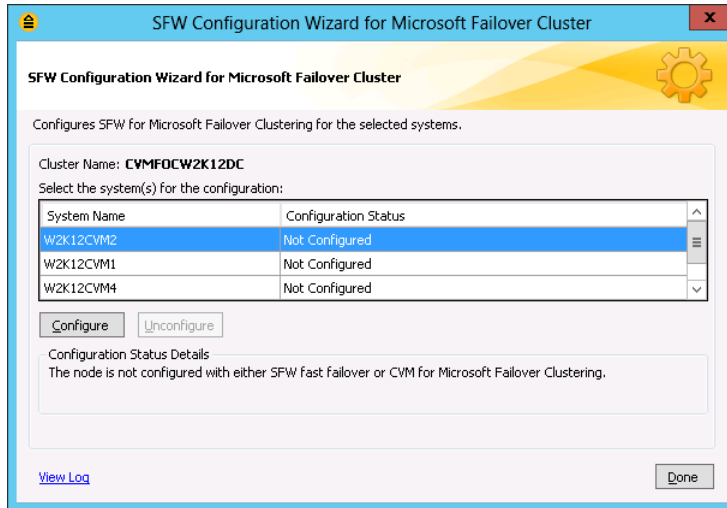
Warning: Please note that this requires application downtime.

- If UAC is enabled, run the program or commands in the “Run as administrator” mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

For information about unconfiguring a CVM cluster, See [“Unconfiguring a CVM cluster”](#) on page 443.

Using the SFW Configuration Wizard for Microsoft Failover Cluster

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), type **Run**, and then press **Enter** to open the **Run** dialog box.
In the **Run** dialog box, type **scc**, and then click **OK** to launch the Solutions Configuration Center (SCC).
- 2 From the Solutions view in the SCC, click on **SFW Configuration Wizard for Microsoft Failover Cluster** to launch the wizard.
- 3 The SFW Configuration Wizard for Microsoft Failover Cluster panel displays cluster and system (node) information and lets you configure or unconfigure systems for a CVM cluster as described below.



Fields

Description

Cluster Name	The automatically-assigned name for the CVM cluster configuration. This name is the same as that of the Microsoft failover cluster.
System Name	This table lists all the systems (nodes) that can participate in a CVM cluster. Select the required cluster nodes that you want to configure for CVM.
Configuration Status	<p>The following configuration statuses are displayed:</p> <ul style="list-style-type: none"> ■ Configured: The node is configured in the CVM cluster. ■ Not Configured: The node is not configured in the CVM cluster. ■ Invalid Configuration: There are configuration issues with the node. ■ Not Supported: Unsupported configuration. This is displayed if you don't select the Microsoft Failover Cluster option during the product installation.
Configuration Status Details	Displays CVM configuration status information for the selected systems.

To configure CVM, select one or more systems that are not yet configured and click **Configure** to configure CVM support for the selected systems.

Note: Ensure that Microsoft failover cluster and CVM membership is same.

Note: When you configure CVM for the first time, even if one system is selected, the wizard configures all the nodes, including those that are not selected.

- 4 Click **Done** to exit the wizard.

Configuring CVM using `ffconfigcli`

This section provides information about using the CLI-based `ffconfigcli` utility to prepare a cluster for CVM. Using the `ffconfigcli` utility from CLI, you can perform the operations that are listed below.

Note: In the following commands, you can use the optional `-verbose` option at the end of a command to display information about the tasks that the command performs. For example, using the `-verbose` option in the following command displays information about the tasks that are performed for auto-configuring a CVM:

```
ffconfigcli -autoconfigure -verbose
```

Note: In the following commands, wherever the name of the node is required, you cannot optionally provide the node's IP address as it is not supported.

For information about unconfiguring a CVM cluster, See [“Unconfiguring a CVM cluster”](#) on page 443.

- **Auto-configuring a CVM cluster:** Run the following command to automatically set up a CVM cluster with the same name as the Microsoft failover cluster:

```
ffconfigcli -autoconfigure
```

You must use this command first if you configure a new CVM cluster.

- **Adding a node to a CVM cluster:** Run the following command to add a node to a configured CVM cluster:

Note: You must use this command only after you configure a CVM cluster using the `ffconfigcli -autoconfigure` command. Also, to successfully add a node to the CVM cluster, the node must be added to the Microsoft failover cluster first.

```
ffconfigcli -addnode <NodeName>
```

Where, `<NodeName>` is the name of the node that you want to add to the CVM cluster. This is a mandatory parameter.

Adding multiple nodes to a CVM cluster: Using the `-addnode` command, you can add only one node at a time. You can add multiple nodes by running the command multiple times in one of the following ways:

- Add the nodes one-by-one to both Microsoft failover cluster and CVM.
For example: First, add node X to Microsoft failover cluster then add it to CVM, add node Y to Microsoft failover cluster first and then to CVM, and so on.
- Add all the nodes to Microsoft failover cluster first and then add them one-by-one to CVM.
For example: Add nodes X, Y, and Z to Microsoft failover cluster first. Then, add the same three nodes to CVM.
However, if you use this method, CVM auto-starts only on the node added last (node Z in the example). For the other previously-added nodes, CVM needs to be restarted. For information about restarting CVM, See [“Restarting CVM”](#) on page 570.

Note: Alternatively, to easily add multiple nodes, you can use the GUI-based SFW Configuration Wizard for Microsoft Failover Cluster.

Creating cluster-shared disk groups and volumes

Once you have run the SFW Configuration Wizard for Microsoft Failover Cluster or its CLI equivalent `ffconfigcli` utility, you need to create a cluster-shared disk group (CSDG) and a cluster-shared volume using the VEA GUI to complete a CVM cluster configuration.

Creating a cluster-shared disk group

Once the cluster-shared disk group (CSDG) is created, you need to create a cluster-shared volume to complete the CVM cluster configuration.

Creating a cluster-shared disk group

- 1 Open and connect to Veritas Enterprise Administrator (VEA).
- 2 In the VEA GUI, either right-click the **Disk Groups** folder in the tree view and then click **New Dynamic Disk Group** or click **New Group** on the toolbar.
- 3 The **New Dynamic Disk Group** Wizard appears. Click **Next** to continue.
- 4 On the next panel, do the following:
 - Type a name for the new disk group in the **Group Name** entry box.

- Select the **Create Cluster Shared Disk Group** option.
- Select the basic disks that you want to include in the new disk group by moving them from the left pane to the right pane by using the **Select** option.

Click **Next** to continue.

- 5** On the next panel, click **Next** to confirm the set of disks to be included in the cluster-shared disk group.
- 6** On the next panel, click **Finish** to create the new cluster-shared disk group.

Once the cluster-shared disk group is created, it appears under the **Disk Groups** folder in the tree view of VEA GUI.

Note: If you receive the error “V-76-58645-20098: Configuration error. Check if GAB/LLT is properly configured”, then you need to configure SFW for Microsoft Failover Clustering using the SFW Configuration Wizard for Microsoft Failover Cluster. If a CVM cluster is already configured, then run the wizard again to find and resolve any configuration-related issues.

Creating a cluster-shared volume

Once the cluster-shared disk group (CSDG) is created, you need to create a cluster-shared volume under the new CSDG to complete the CVM cluster configuration. A dynamic volume that is created under a CSDG is automatically configured in the Microsoft failover cluster as a Volume Manager Shared Volume resource.

Note: You must ensure that the selected drive letter for the new cluster-shared volume is available and not in use on any of the cluster nodes.

Note: In Microsoft failover cluster, you can define the restart attempts on failure of a volume resource under the Policies tab of the resource properties. For this, it is recommended that you do not change the default value of 15 minutes for the “Period for restarts” option. Changing this value to less than 5 minutes causes the volume resource online operation to hang. Moreover, changing the value of the “Maximum restarts in the specified period” option to greater than 5 may create a delay of more than 2 minutes in bringing the volume resource online if any of the earlier restart attempts had failed.

For information about creating a cluster-shared volume, See [“Creating a dynamic volume”](#) on page 48.

Manually creating a Volume Manager Shared Volume resource

When you create a cluster-shared volume, it is automatically configured in the Microsoft failover cluster as a Volume Manager Shared Volume resource.

When a Volume Manager Shared Volume resource is created, it comes online immediately on the node that owns the Microsoft Failover Clustering Available Storage group. You can verify the created resource using the GUI-based Microsoft Failover Cluster Manager snap-in. The Volume Manager Shared Volume resource gets created under **Storage (Storage > Disks** in Windows Server 2012) in Failover Cluster Manager. Once the resource is created and online in the Available Storage group, you can create virtual machines with it or move it to other virtual machines or service groups.

However, the Volume Manager Shared Volume resource is not created automatically in the following cases:

- After upgrading the disk group to the latest version and importing a disk group as a cluster-shared disk group (CSDG).
- If you create a cluster-shared volume using the VEA GUI and choose not to format the volume (by clearing the **Format this volume** check box).
- If you create a cluster-shared volume using the SFW CLI, then the resource is not created automatically because, by default, the command creates a RAW, unformatted volume.

In these cases, you need to manually create a Volume Manager Shared Volume resource for all volumes in a CSDG using Failover Cluster Manager. The following procedure gives information about manually creating the resource for a cluster-shared volume using Failover Cluster Manager.

Note: Before manually creating the resource, you must format the cluster-shared volume with NTFS and ReFS using the VEA GUI and mount it on the node where you want to create the resource.

Manually creating Volume Manager Shared Volume resource in Failover Cluster Manager

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

Launch the Failover Cluster Manager snap-in by clicking **Failover Cluster Manager**.

- 2 On Windows Server 2012, right-click **Roles**. Select **More Actions > Create Empty Role**. This creates a new role.

On Windows Server 2008 and 2008 R2, right-click **Services and Applications**. Select **More Actions > Create Empty Service or Application**. This creates a service group.

- 3 Right-click the new role (called service group in Windows Server 2008 or 2008 R2) and select **Add a Resource > More Resources > Add Volume Manager Shared Volume** from the context menu.

A new Volume Manager Shared Volume resource is created with a default name.

- 4 Right-click the Volume Manager Shared Volume resource and select **Properties**. Complete the following on the **Properties** window:
 - Under the **General** tab, change the default name of the Volume Manager Shared Volume resource to a name of your choice. For example, CSDG1.
 - Under the **Properties** tab, for the **MountPoint** property, provide the mount point of the cluster-shared volume.

Click **Apply** to close the window.

- 5 Right-click the Volume Manager Shared Volume resource and select **Bring this resource online** to bring the resource online on the node with the ownership of the Available Storage group. This also brings the cluster-shared volume online in the VEA GUI.

Configuring a role for volume-level failover

Configuring a role for volume-level failover configures a Hyper-V virtual machine or an application for high availability.

This task involves the following steps:

1. Configure a role for each Hyper-V virtual machine or an application that you want to configure for high availability.

The wizard creates a resource for the configured role.

2. Add storage to the configured role to add a cluster-shared volume as a storage resource
The wizard adds a cluster-shared volume resource to the configured role.
3. For the configured role, set dependency between the two resources

Note: You must perform this step only if the host machine runs Windows 2012 R2 Hyper-V operating system.

These steps should be performed using Failover Cluster Manager. For details refer to Microsoft documentation.

Administering CVM

This section provides information for the following Cluster Volume Manager (CVM) administration tasks:

- Administering CVM using the CLI
See [“Administering CVM using the CLI”](#) on page 571.
- Access modes for cluster-shared volumes
See [“Access modes for cluster-shared volumes”](#) on page 434.
- Commands that are supported for running when the volume is online
See [“Operations supported for running where volume is online”](#) on page 446.
- Storage disconnectivity and CVM disk detach policy
See [“Storage disconnectivity and CVM disk detach policy”](#) on page 437.

Note: If you encounter issues while working with CVM, refer to the troubleshooting section for assistance. See [“CVM issues”](#) on page 562.

Configuring CVM links for multi-subnet cluster networks

This section provides information about using the CLI-based `SFWConnectionManager` utility to configure a CVM link for the nodes on multi-subnet cluster networks. To successfully create a CVM cluster, the participating cluster nodes must share a common cluster network. However, in case of a cross-subnet configuration, where you have nodes across different cluster networks, you need to use the `SFWConnectionManager` utility to create a CVM link to establish cluster and network communications across the nodes that forms a CVM cluster.

You must create or delete a CVM link using `SFWConnectionManager` before you run the SFW Configuration Wizard for Microsoft Failover Cluster (or its CLI equivalent `ffconfigcli`). If you make any changes to the existing CVM link of the CVM cluster using `SFWConnectionManager` after running the wizard (or `ffconfigcli`), then you need to unconfigure and re-configure CVM cluster for the changes to be reflected.

Using the `SFWConnectionManager` utility from CLI, you can perform the operations that are listed below:

- **Creating a CVM link:** Run the following command to create a CVM link across cluster networks:

```
SFWConnectionManager -createlink <ClusterNetworkName  
1>, <ClusterNetworkName 2>, ...<ClusterNetworkName n>
```

This command creates a CVM link for the list of cluster networks provided.

Note: If a cluster network's name has spaces, then it must be specified within double quotes.

- **Deleting a CVM link:** Run the following command to delete an existing CVM link:

```
SFWConnectionManager -deletelink <CVMLinkName>
```

Note: In place of `<CVMLinkName>`, use the `-all` option at the end of the command to delete all the CVM links.

- **Displaying CVM link information:** Run the following command to display information about the CVM links and their cluster networks:

```
SFWConnectionManager -displaylinks
```

Access modes for cluster-shared volumes

The Cluster Volume Manager (CVM) architecture is such that two different volumes can be accessed from two different nodes but same volume cannot be accessed from different nodes. Volumes from the same cluster-shared disk group can be shared exclusively from multiple nodes in a Microsoft failover cluster.

In such a shared storage pool, where any node in the cluster can access a disk group, volume access modes are used to avoid data corruption. Volume access is defined in terms of nodes where only one node in a given cluster is granted access to the volume.

When a volume is brought online on a node, then that node is given access to the volume, while the rest of the nodes cannot access the volume. To give another node access to the same volume, you must first bring the volume offline on the node where it is online, and then bring it online on the other node.

The following states are applicable for cluster-shared volumes:

- **Online (Read/write)**
Allows the data to be read from or written to a shared volume in the storage pool.
- **Offline (Read-only)**
This state means that you cannot access data the volume.

Setting the access mode using the Volume Manager Shared Volume resource

When you create a shared volume on a Cluster Shared Dynamic Disk Group (CSDG) in VEA GUI, the corresponding volume resource automatically gets created under the available storage in a Microsoft failover cluster. The Volume Manager Shared Volume resource gets created under **Storage (Storage > Disks** in Windows Server 2012) in Failover Cluster Manager.

For more information about the resource, See [“Manually creating a Volume Manager Shared Volume resource”](#) on page 431.

Using Failover Cluster Manager, you can bring the Volume Manager Shared Volume resource online or offline, which, in turn, brings the shared volume online or offline.

Setting the access mode using the GUI

You can change access modes for a shared volume that is part of the cluster-shared disk group using the VEA GUI.

From the **Volumes** tab in VEA GUI, right-click the volume and then change the access mode (Online or Offline) as required.

Set the access mode using the CLI

You can change access modes for a shared volume that is part of the cluster-shared disk group using the CLI commands as explained below.

For volume access:

```
vxassist [-g<DynamicDiskGroupName>] <online | offline>  
<VolumeName|DriveLetter|VmName|DrivePath>
```

Brings the volume online or offline on the node. Bring the volume online so that the node is given access to the volume. The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal volume name parameter.)
<code>online</code>	Brings the volume online on the node and gives the node access to it.
<code>offline</code>	Brings the volume offline on the node and the node cannot access it.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder. Must be an empty folder on a local NTFS volume.

Example: To bring the volume online on the node and give the node access to it:

```
vxassist online x:
vxassist -gDG1 online Volume1
```

To check the volume state

Using the GUI:

Using the VEA GUI, you can check the status of a volume on a node (whether it is online or offline). The status is displayed in the State column of the **Volume** tab view. You can also see the volume status by right-clicking the volume and selecting Properties.

Using the CLI:

Using the CLI, you can use the following command to check the status of the volume on a node, whether it is online or offline.

```
vxvol [-g<DynamicDiskGroupName>] volinfo
<VolumeName|DriveLetter|VmName|DrivePath>
```

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal volume name parameter.)
---	--

<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder. Must be an empty folder on a local NTFS volume.

Examples:

```
vxvol -gDG1 volinfo Volume1
```

```
vxvol volinfo i:
```

The "Online" property of the command output provides information about the current status of the volume on the node. The value "Yes" indicates that the volume is online and the value "No" indicates that it is offline.

Storage disconnectivity and CVM disk detach policy

When Cluster Volume Manager (CVM) detects a connectivity failure to one or more plexes on one or more data disks, CVM determines the handling of the connectivity failure based on four types of storage connectivity failures as described below.

See ["Types of storage connectivity failures"](#) on page 437.

See ["How CVM handles local storage disconnectivity with the local detach policy"](#) on page 440.

See ["CVM storage disconnectivity scenarios"](#) on page 440.

See ["Availability of cluster nodes and cluster-shared disk groups"](#) on page 443.

See ["Cluster-shared disk group fail policy"](#) on page 443.

Types of storage connectivity failures


CVM determines the type of storage disconnectivity failure based on the scope of failure. CVM determines whether the failure affects all the nodes (global failure) or only some particular nodes (local failure). CVM also determines whether the failure affects one or more plexes of the volume. If the failure affects all plexes, it is considered as a total failure. Otherwise, it is considered as a partial failure.

CVM defines the following types of storage disconnectivity:

Global partial failure

The following figure illustrates a global partial failure.

Figure 14-1 Global partial failure

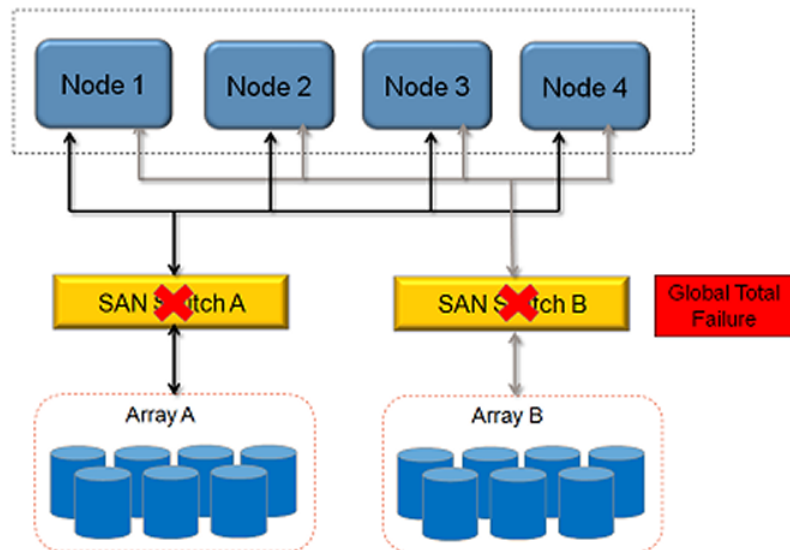


A global partial failure indicates that all of the nodes are affected, but not all of the plexes in the volume. In the above example, all the nodes in the cluster lost access to Array B, which has plex B for the volume.

Global total failure

The following figure illustrates a global total failure.

Figure 14-2 Global total failure




A global total failure means that all the nodes and all plexes in the volume are affected.

Local partial failure

The following figure illustrates a local partial failure.

Figure 14-3 Local partial failure




A local partial failure indicates that the failure did not affect all plexes, and the failure occurred on one or more nodes, but not on all the nodes.

Local total failure

The following figure illustrates a local total failure.

Figure 14-4 Local total failure



A local total failure indicates that the failure affected all the plexes, and the failure occurred on one or more nodes, but not on all nodes.

How CVM handles local storage disconnectivity with the local detach policy

The local detach policy indicates that when a node has an I/O failure to a plex due to a local connectivity issue, the I/Os fail to the local volume. The plex is not detached for the whole cluster. This behavior ensures that all of the plexes are available for I/O on the other nodes. Only the node or nodes that had the failure are affected.

The benefit with this policy is that the redundancy of the volume is protected, as much as possible. The local detach policy supports failover applications in large clusters where the redundancy of the volume is more important than the number of nodes that can access the volume. In other words, you would prefer one or more nodes to lose I/O access to the volume rather than lose a plex of the volume for all of the nodes in the cluster.

Table 14-1 Cluster behavior under I/O failure to a mirrored volume for local disk detach policy

Type of failure	Local detach policy
Global partial failure	Detaches the plex
Global total failure	Disables the volume
Local partial failure	Fails I/O to the volume from the nodes that cannot access the plex
Local total failure	Fails I/O to the volume from the nodes that cannot access the plex

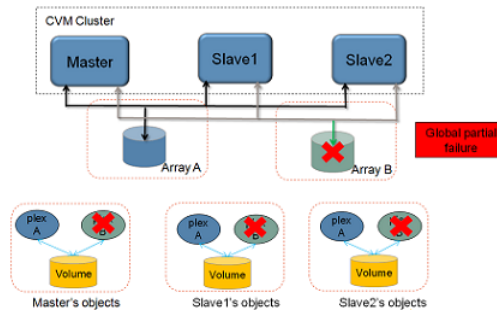
CVM storage disconnectivity scenarios

For the storage disconnectivity failures that are discussed in the previous section, this section provides information about how CVM handles these storage failures using the local detach policy:

Global partial failure

The following figure illustrates how CVM handles a global partial failure with the local detach policy.

Figure 14-5 Local detach policy for global partial failure




A global partial failure indicates that all the nodes are affected, but not all the plexes in the volume. In the above example, all the nodes in the cluster lost access to Array B, which has plex B for the volume. In this case, plex B is detached. None of the nodes can access the plex, so the plex must be detached to maintain the consistency of the mirror. I/O continues to other plexes in the volume. This reduces the redundancy of the volume.

Global total failure

The following figure illustrates how CVM handles a global total failure with the local detach policy.

Figure 14-6 Local detach policy for global total failure

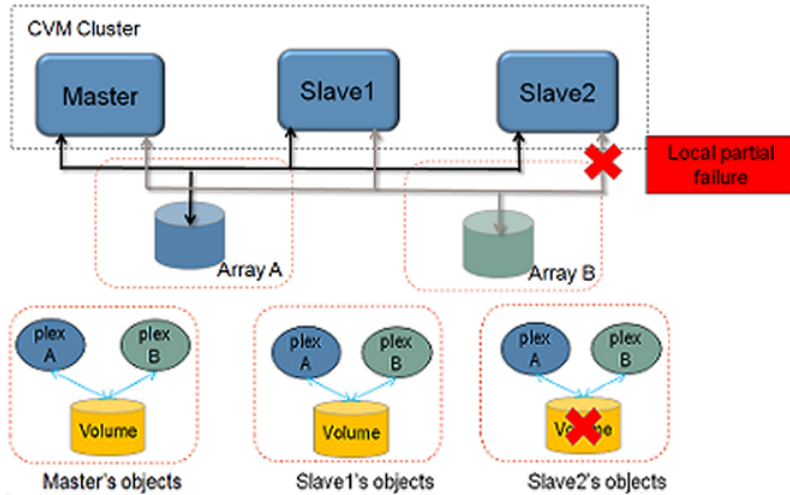


A global total failure indicates that all the nodes and all the plexes in the volume are affected. In this case, the volume is disabled. Since no plexes are available, the volume is not available for any I/Os. If the failure occurs on all the nodes at the same time, no plexes are detached.

Local partial failure

The following figure illustrates how CVM handles a local partial failure with the local detach policy.

Figure 14-7 Local detach policy for local partial failure

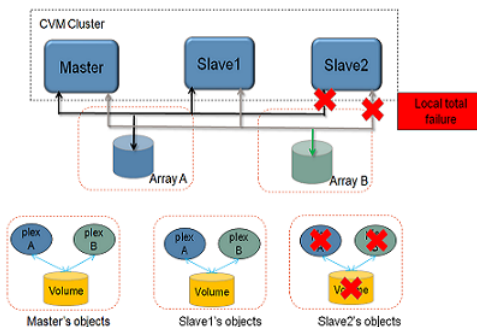


In case of a local partial failure, CVM locally fails to write I/Os to the volume. The local detach policy indicates that CVM should ensure that a local connectivity error only affects the local node. When the I/O fails locally to the volume, applications need to be failed over to another node.

Local total failure

The following figure illustrates how CVM handles a local total failure with the local detach policy.

Figure 14-8 Local detach policy for local total failure



In case of a local total failure for all of the plexes in the volume, CVM locally fails to write I/Os to the volume. I/O fails to the volume from the node or nodes that cannot access the plex. I/O continues to the volume from the other nodes.

Availability of cluster nodes and cluster-shared disk groups

By default, Cluster Volume Manager (CVM) requires that a node joining a cluster must have access to all the disks in all the cluster-shared disk groups (CSDGs). This requirement applies also to a node that was removed from the cluster. The node cannot rejoin the cluster until the disk access is restored.

Similarly, a cluster-shared disk group cannot be imported unless all of the nodes in the cluster can access all of the disks in the disk group.

Cluster-shared disk group fail policy

In the event of loss of connectivity to all the disks of a cluster-shared disk group from the CVM Master node, the Master node leaves the cluster allowing another node to become the Master.

Unconfiguring a CVM cluster

This section provides information about unconfiguring a CVM cluster using the GUI-based SFW Configuration Wizard for Microsoft Failover Cluster and the CLI-based `ffconfigcli.exe` utility. The section provides information about unconfiguring the entire cluster as well as about removing one or more nodes from the CVM cluster.

Please note the following important information before unconfiguring a CVM cluster using the wizard or utility:

- While unconfiguring a CVM cluster, none of the CVM resources should be online on any of the cluster nodes. Similarly, while removing a node, no resources should be online on the node that you want to remove.
- To run the wizard on a node, the CLUSSVC service must be running irrespective of the operation being carried out. However, if you are running the wizard on a node that is not going to be removed, then the CLUSSVC service can be stopped on the node being removed.

On the other hand, the CLI utility functions differently from the wizard. For adding a node and configuring a cluster, the utility requires that the CLUSSVC service is running. However, for removing a node and unconfiguring a cluster, it does not matter if the CLUSSVC service is running or stopped on any node.

- Unconfiguring a CVM cluster does not delete any CVM link information that the wizard (or utility) automatically populates or the administrative user manually creates.

For information about configuring a CVM cluster, See [“Configuring a CVM cluster”](#) on page 425.

Using the SFW Configuration Wizard for Microsoft Failover Cluster

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), type **Run**, and then press **Enter** to open the **Run** dialog box.

In the **Run** dialog box, type **scc**, and then click **OK** to launch the Solutions Configuration Center (SCC).

- 2 From the Solutions view in the SCC, click on **SFW Configuration Wizard for Microsoft Failover Cluster** to launch the wizard.

- 3 The SFW Configuration Wizard for Microsoft Failover Cluster panel displays cluster and system (node) information and lets you configure or unconfigure systems for a CVM cluster.

To unconfigure CVM, select one or more configured systems and click **Unconfigure** to unconfigure CVM support for the selected systems.

Select all the systems to unconfigure the entire CVM cluster.

- 4 Click **Done** to exit the wizard.

Unconfiguring CVM using ffconfigcli

This section provides information about removing a node or unconfiguring a CVM cluster using the CLI-based `ffconfigcli` utility.

Note: In the following commands, you can use the optional `-verbose` option at the end of a command to display information about the tasks that the command performs.

Note: In the following commands, wherever the name of the node is required, you cannot optionally provide the node's IP address as it is not supported.

For information about configuring a CVM cluster, See [“Configuring a CVM cluster”](#) on page 425.

- **Unconfiguring a CVM cluster:** Run the following command to unconfigure and delete a CVM cluster:

Note: It is recommended that you unconfigure the CVM cluster from the Microsoft failover cluster first, and then run this command.

```
ffconfigcli -unconfigure
```

- **Removing a node from a CVM cluster:** After removing a node from the Microsoft failover cluster, run the following command to remove it from the CVM cluster:

```
ffconfigcli -removenode <NodeName>
```

Where, <NodeName> is the name of the node that you want to remove from the CVM cluster. This is a mandatory parameter.

Command shipping

Clustered Volume Manager (CVM) designates one of the cluster nodes as the Master node for the cluster and all the other nodes become Slave nodes. The CVM architecture allows the administrative operations for cluster-shared disk groups (CSDGs), such as disk group creation and volume creation, only on the Master node. With the command shipping feature, the user can execute the supported commands from any of the nodes in the cluster.

The SFW configuration view may differ from node to node in the CVM cluster. For example, disk access (DA) names of the disks in cluster-shared disk groups can differ from node to node. Therefore, you cannot run the commands from the Slave nodes using the SFW configuration view on that node.

Note the following error handling for the commands that you issue from the Slave node, which CVM runs on the Master:

- If the `vxconfigd` daemon on either the Slave node or on the Master node fails, the command exits. The instance of the command on the Master also exits. To determine if the command executed successfully, use the `vxprint` command to check the status of the VxVM objects.
- If the Slave node that shipped the command or the Master node leaves the cluster while the Master executes the command, the command exits on the Master node as well as on the Slave node. To determine if the command executed successfully, use the `vxprint` command to check the status of the VxVM objects.

Operations supported by command shipping

This section provides a list of SFW operations that are supported by the command shipping feature. Therefore, the supported commands for these operations can be run from a Slave node.

Using the command shipping feature, any of these commands that you initiate from a Slave node for a cluster-shared disk group (CSDG), CVM executes them on the Master node automatically.

Table 14-2 Operations supported by command shipping

Operations	Supported Commands
vxassist	vxassist growby vxassist shrinkby vxassist snapback vxassist snapshot
vxcache	vxcache enable vxcache disable
vxdg	vxdg import vxdg join vxdg reclaim vxdg reclaimcancel vxdg split
vxsnap	vxsnap create vxsnap reattach vxsnap restore

Operations supported for running where volume is online

This section provides a list of SFW operations for a cluster-shared disk group (CSDG) that are supported only on the node where the volume is online.

Table 14-3 Operations supported for running where the volume is online

Operations	Supported Commands
vxassist	vxassist growby vxassist shrinkby
vxsnap	vxsnap create vxsnap reattach vxsnap restore
vxvol	vxvol growfs vxvol reclaim vxvol reclaimcancel

About I/O Fencing

Note: I/O Fencing is completely automatic. You do not need to configure anything manually to utilize I/O Fencing.

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

The fencing operation determines the following:

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

This decision prevents possible data corruption.

In the event of a network partition, SFW I/O Fencing technology uses the outcome of Microsoft Failover Clustering arbitration process to choose which nodes to allow access to the shared storage and which nodes to eject from the cluster.

About preventing data corruption with I/O Fencing

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

To provide high availability, Microsoft Failover Clustering takes corrective action when a node fails. In this situation, SFW also updates its components to reflect the altered CVM membership.

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

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

About I/O Fencing operations

I/O Fencing performs identically on node failures and communications failures. When Microsoft Failover Clustering on a node becomes aware of a change in cluster membership, it immediately begins the arbitration process. Depending on the outcome of the arbitration process, the winning node ejects the departed nodes from the data disks. In a split-brain scenario, both sides of the split would race for proving majority. The side that is able to prove the majority, wins the race and fences the losing side.

About I/O Fencing components

The shared storage for SFW must support SCSI-3 persistent reservations to enable I/O Fencing. SFW involves data disks as the shared storage to store share data.

Data disks are standard disk devices for data storage and are either physical disks or Logical Units (LUNs). These disks must support SCSI-3 PR and must be part of standard CVM disk groups. CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

Administering site-aware allocation for campus clusters

This chapter includes the following topics:

- [Overview](#)
- [About the site-aware read policy](#)
- [Administering hosts for site-based allocation](#)
- [Administering disks for site-based allocation](#)
- [Administering volumes for site-based allocation](#)
- [Administering mirrored plexes for site-based allocation](#)

Overview

In a campus cluster or remote mirror configuration, the hosts and storage of a cluster are allocated between two or more sites. These sites are typically connected through a redundant high-capacity network or Fibre Channel that provides access to the storage and communication between the cluster nodes.

If a disk group is configured with storage at multiple sites and if inter-site communication is disrupted, then a serious split-brain condition may occur. This happens when each site continues to update the local disk group configuration copies without being aware of site disruption. For service(s) to come up on a site while other sites are down, a complete (at least one complete plex for each volume) copy of data is needed. Currently, there is no mechanism to ensure that all volumes

have a complete data plex at each site. Data plex keeps getting changed when a volume is resized, a disk is relocated, or when a new volume is added.

Site-aware allocation feature enables applications and services to function properly at a site when other sites become inaccessible. It means that even during site disruption at least one complete plex of a volume is available at each site. Such type of allocation is known as site-based allocation. Users can specify sites when creating volumes or mirrors, and site boundary limits are maintained for operations like volume grow, subdisk move, and disk relocation. Site boundary limit is crossed when a plex is not totally within a site and allocation of the plex crosses the available site boundary.

Site-aware allocation facilitates following types of site-based allocation:

- Site Confined allocation
- Site Separated allocation

The following table describes the terms that are used in the context of site-aware allocation.

Table 15-1 Site-aware allocation and related terminology

Terminology	Description
Site	Logical representation of a set of hosts and set of arrays or enclosures.
Site Separated	<p>Storage for a volume can be taken from a site (s) that are specified during volume creation. Storage from multiple sites is supported for such type of allocation.</p> <p>Storage for a volume is allocated so that each plex of the volume resides completely on the same site, i.e., if a Site Separated volume has two plexes on two sites A and B, each plex resides completely on a separate site.</p> <p>Volume Resize, relocation, relayout, or any such operation keeps each plex on its own site.</p> <p>Multiple plexes can reside on the same site</p>
Site Confined	<p>Storage for a volume can be taken from only a site that is specified while creating a volume. Multiple sites cannot be allocated for such type of volume allocation. The volume resides entirely on the same site.</p> <p>Resize, relocation, relayout, or any such operation only uses storage from the same site.</p>

Table 15-1 Site-aware allocation and related terminology (*continued*)

Terminology	Description
Siteless	<p>Refers to a volume that is not tagged with any site information or site properties. By default, all volumes are Siteless.</p> <p>Note: On upgrading to SFW 7.1 from any previous release versions (which did not have the Siteless option), all volume types are “Siteless” by default. You can manually change the property of volumes after upgrading to either Site Confined or Site Separated provided that conditions like “volume need to be entirely on the same site” or “each plex of the volume resides entirely on a site” are met.</p>
Site boundary	<p>Site boundary limit is said to be crossed when a plex is not totally within a site and allocation of the plex crosses the available site boundary.</p> <p>Automatic operations like hot relocation does not adhere to site boundary restrictions. Storage that is configured with such auto operations become Siteless once site boundary limit is crossed.</p> <p>When a storage becomes siteless, user is notified and Event Viewer logs displays logs to verify the same.</p>

Pre-requisites for enabling site-aware allocation on a campus cluster setup

The following prerequisites are listed below for configuring a campus cluster setup with site aware allocation:

- Each host that is part of a campus cluster setup can be assigned a site name (optional requirement).
You can add a host to a site using the Veritas Enterprise Administrator (VEA) GUI or `vxassist` command line options.
See [“Administering hosts for site-based allocation”](#) on page 452.
- SFW dynamic disks must be tagged with a site name to be able to use the site aware allocation feature. You can add a dynamic disk to a site through the VEA GUI or `vxdisk` command line options.
See [“Administering disks for site-based allocation”](#) on page 454.

About the site-aware read policy

Site-aware read policy is a SFW feature that enhances the I/O read performance. To enhance the read performance, SFW services read I/Os from the plexes at the local site where an application is running. The I/O writes are written to plexes at all sites. By tagging hosts with site information, SFW identifies which hosts belong to which site. The I/O reads initiated by a host from one site are then satisfied by the disks which are tagged with the same site. Tagging hosts and disks with correct site information gives you maximum read performance with the site-aware read policy feature.

Note: The site-aware read policy feature is applicable only for the Site Separated volumes.

The site-aware read policy feature works only if a host is tagged and at least one plex of the site-separated volume resides on disks with the same tags as the host. Also, if a host is tagged, but there are no local site disks available, then the plexes are read sequentially in a "round-robin" fashion. When there is more than one local plex, one of them is chosen as the preferred plex for site-aware read policy.

Note: If any plex is set as a preferred plex to improve I/O performance, then the site-aware read policy feature is not implemented.

If a volume fails over to a site without local plex, then the site-aware read policy becomes the "round-robin" read policy in which the preferred plex that is set before the failover no longer remains applicable.

Administering hosts for site-based allocation

For site-based allocation, hosts can be added or removed from a site through the Veritas Enterprise Administrator (VEA) or command line options.

Note: Site name can take any alphanumeric values from 1 to 9, a...z, A...Z and valid characters like period (.), dash (-), and underscore (_). It can be of maximum 31 characters long. Site names are case insensitive and saved in registry in small case even if user specifies the site names in upper case.

See ["Administering hosts for site-aware allocation through GUI"](#) on page 453.

See ["Administering hosts for site-aware allocation through CLI"](#) on page 453.

Note: If storage is tagged with site information but host is not tagged with any site name, site aware allocation feature can still be used for such storage. However, if the storage is not tagged with any site information, it cannot be used for site-based allocation.

Administering hosts for site-aware allocation through GUI

A host can be tagged with a site name using the Veritas Enterprise Administrator (VEA) GUI.

Adding a host to a site

To add a host to a site through VEA, perform the following steps:

To add a host to a site through GUI

- 1 From the VEA console, select **localhost** and click **Actions > Add Host to Site** from the Menu.

Additionally, you can also right-click **localhost** and select **Site > Add Host to Site** from the context menu.

- 2 On the **Add Host to Site** dialog box, enter the name of the site to which you want to add a host. Say for example, enter the name of site as **site1**. Click **Ok**.

This adds the host to site site1.

Removing a host from a site

To remove a host from a site, perform the following through VEA console:

To remove a host from a site through GUI

- 1 From the VEA GUI, right-click the **local host** and select **Site > Remove Host from Site** option.

Additionally, you can also select Menu option **Actions > Remove Host from Site**.

- 2 A message is displayed to confirm deletion of the host from the site. Click **Yes** to remove a host.

Administering hosts for site-aware allocation through CLI

Use the `vxassist` command options for adding, removing, or displaying the site for a host.

See [“vxassist \(Windows-specific\)”](#) on page 674.

Run the `vxassist taghost site=<sitename>` to add a host to a site named **site1**

See [“vxassist taghost”](#) on page 677.

Use `vxassist showtag` to list the site name to which a host belongs.

See [“vxassist showtag”](#) on page 677.

Use the `vxassist untaghost` to remove the site name for a host.

See [“vxassist untaghost”](#) on page 678.

Administering disks for site-based allocation

For site-based allocation, SFW dynamic disks can be added or removed from a site using the Veritas Enterprise Administrator (VEA) GUI or command line options.

Note: Site name can take any alphanumeric values from 1 to 9, a...z, A...Z and valid characters like period (.), dash (-), and underscore (_). It can be of maximum 31 characters long. Site names are case insensitive and disk is tagged with small case even if user specifies the site names in upper case.

See [“Administering disks for sites through GUI”](#) on page 454.

See [“Administering disks for sites through CLI”](#) on page 456.

See [“View disk properties”](#) on page 110.

Administering disks for sites through GUI

SFW dynamic disks can be added to or removed from the site using the VEA GUI.

Adding a disk to a site

Perform the following steps to add a disk or multiple disks to a site through the VEA console:

To add a disk to a site

- 1 From the VEA console, select a disk and click **Actions > Add Disk to Site** from the Menu.

Additionally, right-click a disk that needs to be added to a site and select the **Add Disk to Site** from the context menu.

Select the site from a list of sites that are seen by a host or specify a new sitename.

- 2 On **Add Disk to a Site** complete the following:

Available Sites

Enable this option if you want to select a site from the drop-down list.

This list contains list of sites to which a host or disk has already been added.

Select a new site

Enable this option if you want to specify a new site name, for example, **site2**.

Available Disks

From this column, select the disk or disks that you want to add to a site. Click **Select >** or **Select All >>** to add the selected disk or disks respectively to a site.

To deselect, click **< Unselect** or **<< Unselect All** option.

Click **OK** to confirm disk selection.

Removing a disk from a site

Note that removing a disk from a site fails when any Site Confined or Site Separated volume resides on the disk.

Perform the following steps to remove a disk from a site through the VEA GUI:

To remove a site from a disk

- 1 From the VEA console, right-click disk or disks that need to be removed from a site. Select **Remove disk from Site** option from the context menu.
- 2 From Remove disk from Site window, select the disk or disks and click **Select >** or **Select All >>** respectively.

Click **OK** to confirm selection and remove the selected disk or disks from a site.

Administering disks for sites through CLI

Disk or disks can be added to or removed from a site using the `vxdisk` command options.

you can use the following `vxdisk` command options:

- To tag a SFW dynamic disk with a sitename run the command:

```
vxdisk setsite site=<sitename> <disklist>
```

where `sitename` is the name of the site and `disklist` is the number of disks that are tagged.
See “[vxdisk setsite](#)” on page 637.
- To display site information or properties of a disk, run the following command:

```
vxdisk diskinfo
```

See “[vxdisk diskinfo](#)” on page 625.
- `vxdisk list`
See “[vxdisk list](#)” on page 628.
- To remove a disk from a site , run the command:

```
vxdisk rmsite
```

Note: `vxdisk rmsite` command will fail on a disk if it has Site Confined or Site Separated volumes resides on the disk.

See “[vxdisk rmsite](#)” on page 637.

Administering volumes for site-based allocation

You can reset or set `Sitetype` property of a SFW dynamic volume either through Veritas Enterprise Administrator (VEA) GUI or command line options.

When performing operations like subdisk move & disk evacuation for site-based volumes, storage must come from the same site as the one to which the volume belongs; otherwise, the operation fails. If you are mirroring a Site-Confined volume or trying to grow a site-based volume (Site-Confined or Site Separated), then also storage should come from the same site as the original volume.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

See [“Expand a dynamic volume”](#) on page 124.

See [“Hot relocation”](#) on page 216.

Setting or resetting volume Sitetype property for site allocation

A new property "Sitetype" has been added to SFW dynamic volumes.

See [“Check partition or volume properties”](#) on page 122.

The Sitetype volume property can take the following values:

- Siteless
- Site Confined
- Site Separated

See [Table 15-1](#) on page 450.

Setting or resetting Volume Sitetype property through GUI

To set or reset the volume property through VEA, perform the following steps:

To set or reset volume property through GUI

- 1 From the VEA console, select a volume from the tree view to select **Actions > Change Volume Sitetype** from Menu.

Additionally, you can also right-click a volume and select **Change Volume Sitetype** from the context menu to set the volume Sitetype property.

Select the following options from the Change Volume Sitetype properties dialog box:

Siteless to Site Separated

All volumes are Siteless by default. You can set the Sitetype property of such volumes to Site Separated.

A Siteless volume is successfully converted to Site Separated provided that each plex of that volume resides entirely on a site.

Siteless to Site Confined

You can reset the Sitetype property of a Siteless volume to Site Confined.

A Siteless volume is successfully converted to Site Confined provided that all the plexes of that volume reside on the same site.

- 2 After selecting the required option, click **OK**.

Additionally, you can also reset the volume Sitetype property of a Site Confined or Site Separated volume to Siteless by right-clicking such a volume and selecting the required option from the context menu.

Setting or resetting volume Sitetype property through CLI

To set or reset volume Sitetype property through the CLI, you can use the following command options:

Use `vxassist set sitetype` command for setting the Sitetype volume property of a Siteless volume to either Site Confined or Site Separated.

```
vxassist -g<DynamicDiskGroupName> set
<VolumeName|DriveLetter|VmName|DrivePath> sitetype={confined |
separated}
```

Use `vxassist resetsitetype` to convert a Site Confined or Site Separated volume to Siteless volume.

```
vxassist -g<DynamicDiskGroupName> resetsitetype
<VolumeName|DriveLetter|VmName|DrivePath>
```

See [“vxassist set sitetype”](#) on page 678.

See [“vxassist resetsitetype”](#) on page 678.

Creating a dynamic volume

For creating site-based volumes (Site Confined or Site Separated), it is a prerequisite that your storage must be site tagged.

See [“Administering disks for site-based allocation”](#) on page 454.

Create SFW dynamic volumes for site-based allocation using the VEA GUI or CLI options.

See [“Creating a volume through GUI”](#) on page 459.

See [“Creating a volume through CLI”](#) on page 459.

Creating a volume through GUI

For information about creating a site-aware volume with site properties through the Veritas Enterprise Administrator (VEA) console or GUI, See [“Creating a dynamic volume”](#) on page 48.

Creating a volume through CLI

You can use the `vxassist make` command to create a volume for site-based allocation.

```
vxassist [-b] [-o notrackalign] -g<DynamicDiskGroupName> make  
[<VolumeName>] <length>[type={mirror|raid5|stripe}[,log]]  
[<alloc_attributes...>] [Column=N] [StripeUnit=<width>]  
[Site=<sitename1>[,<sitename2>,...]]  
[DriveLetter={A|B|..|Z}|DrivePath=<DrivePath>] [<[!]diskname |  
p#c#t#l#> ...]
```

See [“vxassist make”](#) on page 639.

For example run the following command:

```
vxassist -g testdg make testvoll 100M Site=s1
```

The above command creates a Site Confined volume (`testvoll`) on disks belonging to site (`s1`)

For example, run the following command:

```
vxassist -g testdg make testvoll 100M type=mirror mirror=3 Site=s1
```

The above command creates a Site Confined volume (`testvoll1`) with all three plexes residing on disks belonging to site (`s1`.)

For example, run the following command:

```
vxassist -g testdg make testvoll1 100M type=mirror mirror=3  
Site=s1,s2,s3
```

The above command creates a Site Separated volume (`testvoll1`) with each of the three plexes residing on disks belonging to sites (`s1`, `s2`, and `s3`) respectively.

Administering mirrored plexes for site-based allocation

If you create more than one mirror at a time, you may see inconsistent information on the progress bar. Also, the generation of multiple mirrors does affect system resources. After creating a mirror, you may want to wait until it has finished generating before creating another mirror.

See [“Dynamic volume types”](#) on page 56.

Note: Adding mirrors to a volume involves a certain amount of time for mirror resynchronization. Take care not to perform actions on the volume until the command is complete. Also, if you inadvertently shut down the server or deport the disk group containing the volume before mirror resynchronization completes, any mirrors that were being added to the volume are deleted when the disk group with the volume comes online again.

Adding a mirror to a site-based volume through GUI

If you want to add a mirrored plex to a disk or storage that is not site tagged, then you should change the volume `Sitetype` property to `Siteless`.

In case of a Site Confined volume, mirrors can be added within the same site only, say for example **Site1**. If you want to add a mirrored plex to a different site, say **Site2**, then you should change the volume `Sitetype` property to `Site Separated` or `Siteless`.

In case of a Site Separated volume, mirrors can be added to a different site. All volumes that are required for an application to come up must have a complete plex at all the specified sites. Writes are written to the specified sites as done in case of a normal mirrored volume.

See [“Setting or resetting volume `Sitetype` property for site allocation”](#) on page 457.

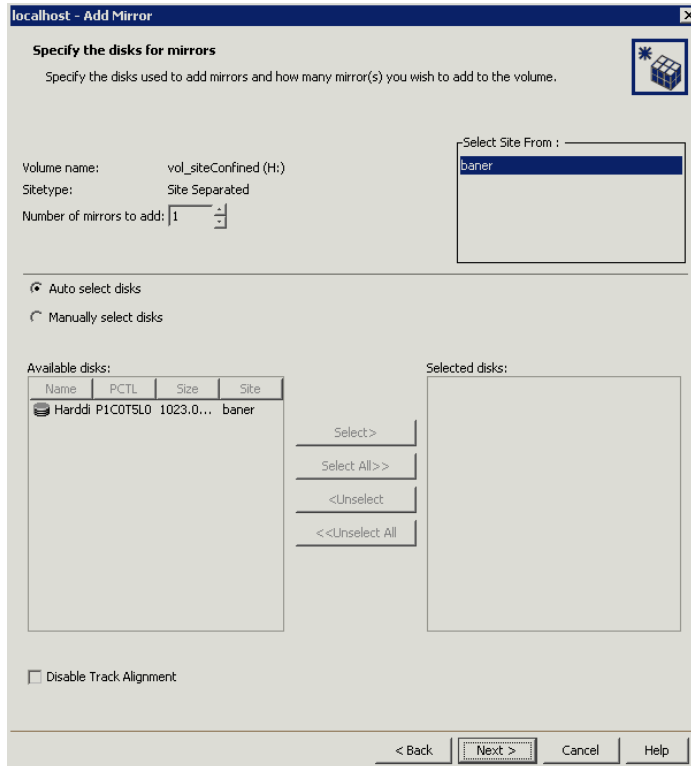
Perform the following steps to add a mirror:

To add a mirror for a site-based volume

- 1** Right-click the volume you want to mirror.

2 Select **Mirror > Add** from the context menu.

The Add Mirror wizard appears. Specify how many mirrors you want to add to the volume, and, if required, manually assign the destination disks for the mirrors.



Complete the following on this screen:

Name of the Volume

Default name of the volume to which you want to add a mirror is displayed

Number of mirrors to add

Select the number of mirrors you want to add from this drop-down list.

Select site from	<p>If the volume is site Separated, then this list is displayed. Select sites as per your requirement. Disks are displayed in Available disks as per specified sites.</p> <p>If volume is Siteless, then this list is grayed out and all available disks are displayed for selection.</p> <p>If the volume is Site Confined, then this site selection list is grayed out and disks from the same site as volume are displayed.</p>
Auto select disks	Enable this option, if you want to assign disks automatically.
Manually select disks	<p>Enable this option, if you want to manually assign disks.</p> <p>The volume site type, whether Site Confined or Site Separated, is displayed based on system auto-selection.</p> <p>Select disk or disks from the Available disk list and click Select or Select All respectively.</p>

Click **Next**.

- 3** In the **Specify the attributes** dialog box, select the Concatenated or Striped volume layout type. If you create a striped mirror, the Columns and Stripe unit size boxes need to have entries. Defaults are provided.

You may also specify a mirror that mirrors across disks with the following attributes:

- Connected by port
- Identified by target
- Contained in an enclosure
- Connected by channel

In addition for a striped layout, you may specify to have the stripe across disks done by the following:

- Port
- Target
- Enclosure

- Channel

Click **Next** to continue.

The operation to create a new dynamic volume fails if the appropriate resources are not available to support the selected attributes to mirror across disks.

- 4 On the summary page, click **Finish** to add the mirror.

Adding a mirror to a site-based volume through CLI

Use the `vxassist mirror` command to add a mirror to an existing site-based volume.

```
vxassist [-b] [-o notrackalign] [-g<DynamicDiskGroupName>] mirror  
<VolumeName|DriveLetter|VmName|DrivePath> [type=stripe]  
[<alloc_attributes ...>] [Column=N] [StripeUnit=<width>]  
[Site=<sitename1>[,<sitename2>,...]] [<[!]diskname | p#c#t#l#> ...]
```

See “[vxassist mirror](#)” on page 650.

For example, run the following command:

```
vxassist -g testdg mirror testvol Site=s1
```

Here, `testvol` is a Site Confined volume residing on-site `s1`. This creates a mirror of volume (`testvol`) on disks belonging to site (`s1`).

For example, run the following command:

```
vxassist -g testdg mirror testvol mirror=2 Site=s2,s3
```

Here, `testvol` is a Site Separated volume residing on-site `s1`. This creates two mirrors of volume (`testvol`) with one plex on disks belonging to site (`s2`) and other on disks belonging to site (`s3`).

SFW for Hyper-V virtual machines

This chapter includes the following topics:

- [Introduction to Storage Foundation solutions for Hyper-V environments](#)
- [Live migration support for SFW dynamic disk group](#)
- [Administering storage migration for SFW and Hyper-V virtual machine volumes](#)
- [Optional Storage Foundation features for Hyper-V environments](#)

Introduction to Storage Foundation solutions for Hyper-V environments

This section includes the following topics:

- See [“Advantages of running Storage Foundation in the Hyper-V parent”](#) on page 466.
- See [“How Storage Foundation manages storage for virtual machines”](#) on page 467.
- See [“Related documentation on Storage Foundation”](#) on page 468.

Advantages of running Storage Foundation in the Hyper-V parent

Storage Foundation (SFW) is a host-level volume manager that provides a means to virtualize the storage that is seen by the host it runs on. SFW provides central-point control of that storage space.

By running SFW in the Microsoft Hyper-V parent partition, SFW features and functionality extend to the virtual machines (VMs), offering benefits that would otherwise be unavailable at the guest level.

See [“How Storage Foundation manages storage for virtual machines”](#) on page 467.

SFW has added the following features and functionality specifically to support Microsoft Hyper-V VMs when running SFW in the Hyper-V parent:

- SFW live migration support
 You can configure the SFW storage on which the VMs reside to support VM live migration between nodes of a Microsoft failover cluster. Detailed instructions are available on how to implement live migration for VMs on SFW storage.
 See [“About implementing Hyper-V virtual machine live migration on SFW storage”](#) on page 470.
- SFW storage migration for VMs
 The SFW storage migration feature enables you to view and select VMs to migrate to different storage. Detailed instructions are available on how to implement the storage migration solution for VMs.
 See [“About storage migration”](#) on page 478.

SFW also offers advanced features and functionality, such as multi-pathing, replication, and snapshots, which further extend the capabilities of Windows in the data center. More information is available on how to use the following features and the benefits they provide in a Hyper-V environment:

- Using Dynamic Multi-Pathing (DMP) to provide failover and load-balancing to the LUNs that host the VMs in the child partition (DMP Device Specific Modules (DSMs) option)
- Replicating VMs between sites (Volume Replicator option)
- Maintaining Quick Recovery snapshots of the VMs (FlashSnap option)

See [“About using optional Storage Foundation features in the Hyper-V parent”](#) on page 491.

How Storage Foundation manages storage for virtual machines

In virtual environments, managing the storage that is used by guests is not an easy task. Typically, the guest is separated from the physical storage. Storage Foundation (SFW) provides several solutions to make it easier to manage storage requirements for virtual machines.

With Hyper-V, guests reside on virtual hard disk (VHD) files, which in turn are located on the volumes that reside on physical storage. Direct access to those volumes or the LUNs they reside on is not available from the guest. The VHD files are provisioned by the parent on storage that is accessed by the parent partition. As storage needs change in the guest VHDs, they may require additional space. It can

be difficult to effectively manage space requirements or to relocate a guest from one storage location to another.

Running SFW in the parent provides the following storage management solutions for VHDs:

- The SFW storage migration feature enables you to view and select VMs to migrate to different storage.
 For details on using SFW for migrating VM to new storage, see the following: See [“About storage migration”](#) on page 478.
- SFW allows for dynamically growing the volumes that host the guest VHDs. As SFW allows for growth of all volume types, the volumes that host the VHD files can be configured for performance by RAID-5, striping or mirrored-stripes.
- In environments that use thin provisioned storage, SFW can be configured to automatically grow volumes based on user- defined space thresholds and policies that set the amount to grow the volumes by and whether that growth should be restricted or unrestricted. This counters the effects of NTFS uncontrolled growth tendencies in a thin environment, by allowing the creation of small volumes on the thin storage, which grows automatically as needed, triggering corresponding growth in the hardware.
- As a host-level volume manager, SFW also allows for mirroring volumes across arrays and, with its support for dynamic disk operations in a cluster, the creation of stretch or campus clusters.

Related documentation on Storage Foundation

The following table lists documentation on SFW that is related to deploying Microsoft Hyper-V virtual machines with SFW.

Table 16-1 Related documentation on SFW solutions

For information about	Refer to
Installation information	<i>Veritas InfoScale Installation and Upgrade Guide</i>
Information on all SFW features	<i>Storage Foundation Administrator’s Guide</i>
Information on Volume Replicator (Volume Replicator)	<i>Volume Replicator Administrator’s Guide</i>

Table 16-1 Related documentation on SFW solutions (*continued*)

For information about	Refer to
Implementing SFW snapshot solutions for applications	SFW Quick Recovery solutions guides. Guides are available for: <ul style="list-style-type: none"> ■ Microsoft Exchange 2007 and 2010 ■ Microsoft SQL Server 2008, 2008 R2, 2012, 2012 R2 ■ Enterprise Vault
Implementing Microsoft Cluster solutions with SFW for applications on physical computers	SFW Microsoft Clustering solutions guides. Guides are available for: <ul style="list-style-type: none"> ■ Microsoft Exchange 2007 ■ Microsoft SQL Server 2008, 2008 R2, 2012, 2012 R2 In addition, the following guide provides general guidelines for using Microsoft Clustering with SFW storage for other applications or server roles: <i>Storage Foundation and High Availability Solutions Solutions Guide</i>

Live migration support for SFW dynamic disk group

This section includes the following topics:

- See [“Tasks for deploying live migration support for Hyper-V virtual machines”](#) on page 470.
- See [“About implementing Hyper-V virtual machine live migration on SFW storage”](#) on page 470.
- See [“Installing Windows Server”](#) on page 472.
- See [“Preparing the host machines”](#) on page 472.
- See [“Installing the SFW option for Microsoft failover cluster option”](#) on page 473.
- See [“Using the SFW Configuration Wizard for Microsoft Failover Cluster for Hyper-V live migration support”](#) on page 473.
- See [“Configuring the SFW storage”](#) on page 474.
- See [“Creating a virtual machine service group”](#) on page 475.

- See [“Setting the dependency of the virtual machine on the VMDg resource”](#) on page 476.

About implementing Hyper-V virtual machine live migration on SFW storage

This chapter provides steps for configuring SFW Microsoft Hyper-V virtual machine live migration on a highly available failover cluster disk group resource.

Live migration of a Hyper-V virtual machine (VM) is achieved through the use of Microsoft Failover Clustering feature. Live migration significantly increases availability of the virtual machines during planned and unplanned downtime.

Live migration produces significantly less downtime for the virtual machines that are migrated. Users can have uninterrupted access to the migrating virtual machine. The guest operating system in the migrating virtual machine is unaware about the ongoing migration. In addition, physical host maintenance can be carried out with no effect on virtual machine availability; this maintenance can occur during normal business hours.

You can configure the SFW storage on which the virtual machine (VM) resides to support VM live migration between nodes of a Microsoft failover cluster.

Note: If you upgrade SFW, and there is a disk group resource already configured in the cluster, then after the upgrade you must take the disk group resource offline and then restart the cluster service. **FastFailOver** attribute for the disk group resource is displayed on the Failover Cluster Manager snap-in. To successfully set it to **True**, configure SFW Fast Failover Configurations for Live Migration support which is between the cluster nodes, using the SFW Configuration Wizard for Microsoft Failover Cluster.

Tasks for deploying live migration support for Hyper-V virtual machines

To deploy SFW Microsoft Hyper-V virtual machine live migration on the host side, perform the following tasks in the sequence shown.

Table 16-2 Process for configuring SFW Hyper-V virtual machine live migration

Action	Description
Review the requirements	See “About implementing Hyper-V virtual machine live migration on SFW storage” on page 470.
Install Windows Server	See “Installing Windows Server” on page 472.
Prepare the host machines	<ul style="list-style-type: none"> ■ Adding the Hyper-V role See “Adding the Hyper-V role” on page 472. ■ Adding the Microsoft Failover Clustering feature on the host side See “Adding the Microsoft Failover Clustering feature” on page 472. ■ Configuring failover cluster nodes on the host side See “Configuring failover cluster nodes” on page 472.
Install InfoScale Storage (select Microsoft Failover Cluster option during the installation)	See “Installing the SFW option for Microsoft failover cluster option” on page 473.
Configure SFW support for Hyper-V live migration using the SFW Configuration Wizard for Microsoft Failover Cluster	See “Using the SFW Configuration Wizard for Microsoft Failover Cluster for Hyper-V live migration support” on page 473.
Configure the storage	<p>Use the Veritas Enterprise Administrator (VEA) console to create disk groups and volumes.</p> <ul style="list-style-type: none"> ■ Creating dynamic cluster disk groups See “Creating dynamic cluster disk groups” on page 474. ■ Creating dynamic volumes See “Creating dynamic volumes” on page 475.
Create a virtual machine service group	See “Creating a virtual machine service group” on page 475.
Set the dependency of the virtual machine on the VMDg resource	See “Setting the dependency of the virtual machine on the VMDg resource” on page 476.

Installing Windows Server

Install Windows Server on the desired host systems. Refer to Microsoft documentation for details.

Preparing the host machines

For virtual machine live migration support, you need to add the Hyper-V role and configure a Microsoft failover cluster on your host machines. Perform the following tasks in the order shown:

Adding the Hyper-V role

After installing Windows Server, the next step is to add the Hyper-V role to enable the live migration feature. To enable the Hyper-V role, refer to Microsoft Hyper-V documentation for details.

Adding the Microsoft Failover Clustering feature

Install the Microsoft Failover Clustering feature on all host systems by using the Add Features option from the Failover Cluster Manager snap-in. Refer to Microsoft documentation for details on installing the Failover Clustering feature.

If you install the Microsoft Failover Clustering feature on a server on which InfoScale Storage is already installed, manually restart the Veritas Enterprise Administrator Service (VxSvc) by running the following CLI commands:

- `net stop vxsvc`
- `net start vxsvc`

Configuring failover cluster nodes

Configure a Microsoft failover cluster on all of your host nodes. Refer to Microsoft documentation for details on how to add a failover cluster node. Ensure that you have fulfilled the clustering network requirements before you start creating the failover cluster nodes.

Note: Verify that you have at least three network adapters (two NICs exclusively for the private network and one for the public network). However, when using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.

Installing the SFW option for Microsoft failover cluster option

Administrative privileges are required on both host and guest operating systems for installing the Storage Foundation Cluster Option for Microsoft failover cluster. The SFW option for Microsoft failover cluster is installed when you install InfoScale Storage and select the Microsoft Failover Cluster option during the installation.

Before you install InfoScale Storage, you must install the Microsoft Failover Clustering feature and configure failover cluster nodes on all the systems that are part of the live migration configuration.

Installing InfoScale Storage requires a restart, but a restart on the active cluster node causes it to fail over. Hence, it is advisable to use a "rolling install" procedure to install InfoScale Storage first on the inactive cluster node. Then, move the cluster resources to the other node and install on the inactive node.

During InfoScale Storage installation using the product installer, make the following selections:

- Select InfoScale Storage as the product to install.
- When selecting the available options from the server components, ensure that you select the **Microsoft failover cluster** option.
 - Leave the client components selected (the default).

For details on using the product installer or command line installation, see the *Veritas InfoScale Installation and Upgrade Guide*.

Using the SFW Configuration Wizard for Microsoft Failover Cluster for Hyper-V live migration support

After configuring the cluster, use the SFW Configuration Wizard for Microsoft Failover Cluster to set up the SFW components for enabling support for live migration. SFW support for Hyper-V live migration is supported on a network within a Microsoft failover cluster. This wizard also enables fast failover support for SFW disk group resources in the cluster.

Enabling Hyper-V live migration support in a Microsoft Failover Clustering environment involves the following tasks:

- Configuring SFW for Hyper-V live migration support using the SFW Configuration Wizard for Microsoft Failover Cluster
- Enabling FastFailover attribute for SFW VMDg resources

You can launch SFW Configuration Wizard for Microsoft Failover Cluster from the Solutions tab in the Solutions Configuration Center (SCC). The wizard can be run from any node of the Microsoft failover cluster that has Hyper-V role added to it.

You can use the wizard for the following scenarios:

- After configuring a Microsoft failover cluster for the first time or after adding a new node to the Microsoft failover cluster
- For enabling FastFailover attribute for SFW VMDg resources
- After removing a node from a Microsoft failover cluster
- Reconfiguring live migration support in case of any change in the network properties

For detailed information about using the wizard to enable Hyper-V live migration support and setting FastFailover attribute for the SFW VMDg resource, See [“Enabling fast failover for disk groups in a Microsoft Failover Clustering environment”](#) on page 172.

For more related, useful information on this, See [“Preparing SFW for Microsoft Failover Clustering”](#) on page 426.

Note: In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS and ReFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode.

Configuring the SFW storage

Using Storage Foundation, you need to create dynamic cluster disk groups and volumes for a cluster environment as described below.

Creating dynamic cluster disk groups

You need to create a dynamic cluster disk group with volumes on shared storage so that they can be shared between nodes in the cluster. Part of the process of creating a dynamic disk group is assigning it a name. You must choose a name that is unique to your environment. Make note of this name, as it is required later.

Note: Setting up a Microsoft failover cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

To create dynamic cluster disk groups, use the Veritas Enterprise Administrator (VEA). The VEA can be invoked on one of the servers and can be used to connect to all the other servers. However, VEA can also be launched on client system and can be used to manage all the servers remotely.

Note: For SFW Hyper-V live migration support, only one virtual machine (VM) per disk group is a mandatory requirement.

While creating a new disk group, you need to select the "Create cluster disk group" option. For information about creating a disk group, See "[Creating a dynamic disk group](#)" on page 44.

Creating dynamic volumes

After creating a dynamic cluster disk group, you need to create a volume for it. For information about creating a volume, See "[Creating a dynamic volume](#)" on page 48.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

Creating a virtual machine service group

After adding a Volume Manager Disk Group (VMDg) resource, proceed with adding a virtual machine on the active failover cluster node.

Note: Virtual machine and virtual hard disk (.vhd) must be stored in the VMDg resource. This is required to make the virtual machine highly available.

To make the shared storage, i.e, the VMDg resource, available to the virtual machine, you must create the virtual machine on a cluster node that owns the storage. You can either create a new virtual hard disk (.vhd), use an existing .vhd as shown in the procedure below, or you can create it later.

To create a virtual machine

- 1 Click **Start > Administrative Tools > Failover Cluster Manager**.

If you are not connected to the cluster node that owns the shared storage connect to it.

- 2 Click on **Service and Applications > Virtual Machine > New > Virtual Machine** from the left pane of the Failover Cluster Manager snap-in.

- 3 The New Virtual Machine Wizard is launched. Review the information on the welcome page.
Click **Next**.
- 4 On the Specify Name and Location page, specify a name for the virtual machine, for example, **VM1**.
- 5 Enable the **Store the virtual machine in a different location** check box, and then type the full path. You can also click **Browse** and copy the virtual hard disk (VHD) file to the Volume Manager Disk Group (VMDg1) resource created in earlier section of this document for storing the virtual machine.
- 6 On the **Memory** page, specify the amount of memory that is required for the operating system that will run on this virtual machine.
- 7 On **Configure Networking** page, select the network adapter that is configured to use for Hyper-V.
- 8 On the **Connect Virtual Hard Disk** page, three options are shown:
 - **Create a new virtual hard disk**
 - **Use existing virtual hard drive**
 - **Create virtual hard disk later**Select the required option and give the path to the VHD (.vhd) file on the Volume Manager Disk Group (VMDg) resource.
The VHD (.vhd) file must be stored in a location where the VMDg resource is located.
- 9 Click **Finish** to close the wizard.

Setting the dependency of the virtual machine on the VMDg resource

Perform this task only if your host machine runs Windows 2012R2 Hyper-V operating system.

As the virtual machine configuration and Virtual Hard Disk (VHD) file lie on the VMDg resource, you need to set a dependency of the virtual machine resource on the VMDg resource to make it highly available and live migrate it.

Perform the following steps to set a dependency of the VMDg resource on the virtual machine:

To set the dependency of the VMDg resource on the virtual machine

- 1 Right click the VMDg resource from the center pane of the Failover Cluster Manager snap-in. Select **Actions > More Actions > Move this resource to another service group or application** from the menu.

VMDg resource is now moved to the selected virtual machine.

- 2 From the center pane of console, select and expand the virtual machine resource to display the **Virtual Machine Configuration Resource**.

Right-click and select **Properties > Dependency** tab to make this Virtual Machine Configuration Resource dependent on the Volume Manager Disk Group resource (VMDg).

Select the VMDg resource that is displayed and click **Ok** to set the dependency.

Now, the virtual machine is ready to be live migrated to other cluster node.

Refer to Microsoft documentation for details regarding live migrating a virtual machine.

Administering storage migration for SFW and Hyper-V virtual machine volumes

This section includes the following topics:

- See [“About storage migration”](#) on page 478.
- See [“About performance tunables for storage migration”](#) on page 478.
- See [“Setting performance tunables for storage migration”](#) on page 478.
- See [“ About performing online storage migration”](#) on page 480.
- See [“Storage migration limitations”](#) on page 481.
- See [“About changing the layout while performing volume migration”](#) on page 481.
- See [“Migrating volumes belonging to SFW dynamic disk groups”](#) on page 481.
- See [“Migrating volumes belonging to Hyper-V virtual machines”](#) on page 483.
- See [“Migrating data from SFW dynamic disks of one enclosure to another”](#) on page 486.
- See [“ Converting your existing Hyper-V configuration to live migration supported configuration”](#) on page 488.

About storage migration

SFW provides the ability to move volumes to new storage locations using the Storage Migration Wizard. The storage migration feature facilitates moving multiple volumes to different set of disks while the volumes are still online. Volumes that are associated with a Hyper-V virtual machine (VM) or an SFW disk group can be moved in a single administrative operation while the volumes are online without stopping the applications or Hyper-V VMs.

Storage migration provides administrators great flexibility when deploying new arrays or moving to LUNs that are configured for better performance.

Note: Volume layout and site boundary constraints are preserved during storage migration.

See [“Administering volumes for site-based allocation”](#) on page 456.

About performance tunables for storage migration

Performance tunables are provided to optimize the time that is taken for the storage migration operation. Using them, you can enhance the storage migration performance, in migrating data from source disks to target disks. These performance tunables give you the flexibility to trade between storage migration performance and application performance.

Usage of performance tunables is limited to storage migration operation only.

The following new performance tunables are introduced:

- I/O size
- Task Threads

You can access these tunables from **VEA > Control Panel > Task Settings > Configuration Task Performance Tuning** tab.

For more details about these tunables,

See [“Setting performance tunables for storage migration”](#) on page 478.

Setting performance tunables for storage migration

The performance tunables are set at default values. However you can define these values, to optimize the storage migration performance:

Administering storage migration for SFW and Hyper-V virtual machine volumes


I/O size	<p>To read or write data to the disks during migration.</p> <p>To trade off between application performance and storage migration performance.</p> <p>The I/O size tuning is not applicable for migration of striped volumes, if you do not change the layout.</p> <p>IO size, default value (in sectors) = 4096, Lower Bound = 128, Upper Bound = 16384</p>
Task threads	<p>To perform the subdisk move</p> <p>Can be specified separately for storage migration tasks</p> <p>Operations that are initiated from array migration and storage migration wizards, use task threads value that is specified under storage migration tunables.</p> <p>Task threads, default value = 1, Lower Bound = 1, Upper Bound = 20</p>
Task Throttling	<p>No new parameters are added to the task throttling tab. However, to further optimize the application performance during storage migration, you can modify the parameters for MirrorCreate (in case of volume relay layout) and SDMove (without volume relay layout and array migration). These parameters are available on the Task Throttling tab.</p>

Perform the following steps from the system where the storage (disk group volumes or array enclosure disks) is mounted and online.

To set the storage migration performance tunables

- 1 Launch VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Connect to the local system.
- 2 Click **Control Panel** in the left panel and then double-click **Task Settings** in the right pane.

- 3 On the localhost - Task Settings window, click the **Configuration task performance tuning** tab to view the new performance tunables.



- 4 In the Storage migration tunable box, define appropriate values for the new tunables, I/O size and Task threads, and then click **OK**.

The attribute values take effect immediately and apply to all the storage migration operations run subsequently.

These settings do not apply to the storage migration operations that are already in progress.

About performing online storage migration

The storage migration wizard migrates data on the disks from one enclosure to another, by moving volumes and optionally changing their layout. This is achieved in a single administrative operation, while the volumes are online, and without stopping the applications.

The following operations are available to perform online storage migration:

- Migrate volumes belonging to SFW dynamic disk groups.
See [“Migrating volumes belonging to SFW dynamic disk groups”](#) on page 481.
- Migrate volumes belonging to Hyper-V virtual machines.
See [“Migrating volumes belonging to Hyper-V virtual machines”](#) on page 483.

- Migrate data from SFW dynamic disks of one enclosure to another. See [“Migrating data from SFW dynamic disks of one enclosure to another”](#) on page 486.

Storage migration limitations

The limitations while performing storage migration are as follows:

- For migrating volumes, the target disks that are selected for migration must belong to the same disk group.
- If the Hyper-V virtual machine configuration spans across multiple disk groups, then storage migration should be performed on a per disk group basis.
- RAID-5 volumes are not supported for migration.

About changing the layout while performing volume migration

You can change the volume layout while migrating it, from concatenated to striped or vice versa. In addition, you can also change the number of columns or stripe unit size of a stripe volume.

For a mirrored volume, the specified target layout applies to all the plexes of that particular volume, except the snap ready plexes.

Migrating volumes belonging to SFW dynamic disk groups

Use this procedure to perform storage migration for SFW dynamic disk groups.

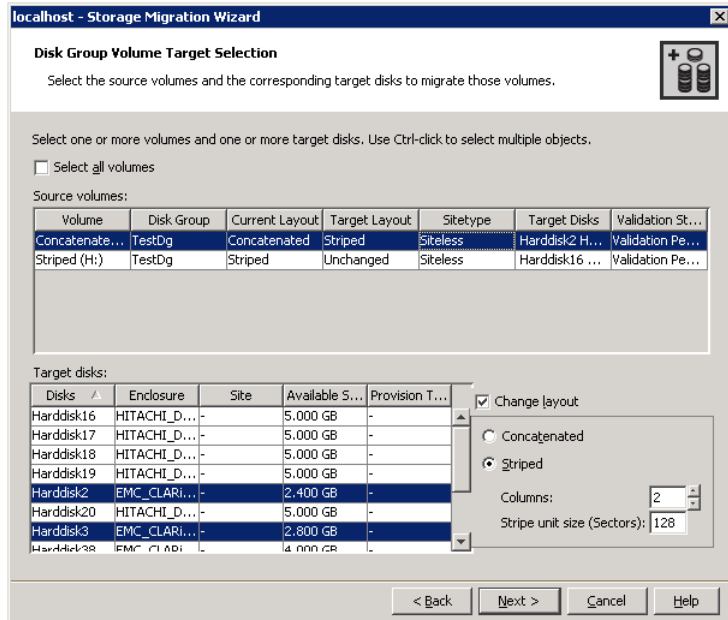
Ensure that the disk groups that you want to migrate are imported on the system where you run this procedure.

To migrate volumes belonging to SFW dynamic disk groups

- 1 Launch VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Connect to the local system.
- 2 In the left panel expand **Disk Groups** and then either right-click and select **Migrate Storage** or from the menu bar click **Actions > Migrate Storage** to start the Storage Migration Wizard.
- 3 Review the prerequisites on the **Welcome** panel and then click **Next**.
- 4 On the **Storage Component Selection** panel, Select **Migrate disk group volumes**, and then select the desired disk group from the drop-down list and click **Next**.

Administering storage migration for SFW and Hyper-V virtual machine volumes

- 5 On the **Disk Group Volume Selection** panel, you can select one or more volumes for online migration. The **Select all volumes** option is selected by default. Clear the **Select all volumes** option and press Ctrl+click if you want to select one or more volumes for online migration. Click **Next**.



- 6 On the **Disk Group Volume Target Selection** panel, you can select the source volumes and the corresponding target disks to migrate volumes. The **Select all volumes** check box is selected by default. Clear the **Select all volumes** check box to map target disks to individual volumes and **Change layout** if desired. Press Ctrl+click to select multiple target disks.
- All selected source volumes can migrate to any target disks if the **Select all volumes** option is selected.
 - If you select the **Select all volumes** option, you do not have the option to change the layout. The layout can be changed only on selecting individual volumes.
 - If you select the **Change layout** check box and opt to change the layout of the volume, the target layout column gets populated with the selected layout, otherwise **Unchanged** is displayed.

Click **Next** to go to the **Summary** panel.

Unless the validation status of all the volumes shows **Success**, the GUI does not navigate to the next panel. The validation status column is populated when you click **Next** on the **Disk Group Volume Target Selection** panel.

- 7 On the **Summary** panel, the mapping of the source volumes and target disks is shown based on the selection that is done on the previous panel. The **Storage migration summary** panel displays the name of the disk group, the source volumes, target layout and the target disks.

If the **select all** option is selected then the summary is shown in the form of a grouping of source volumes and target disks. If the **select all** option is unchecked, the individual selected volumes and assigned target disks is displayed on the **Summary** panel.

The **Save Summary** option on the **Summary** panel lets you save user selection summary in an HTML file.

The default summary file name is

```
StorageMigration_Summary_yyyy_MM_dd_HH_mm_ss.html.
```

The default location of the file is %allusersprofile%\veritas.

- 8 Click **Finish** to begin the volume migration.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

Migrating volumes belonging to Hyper-V virtual machines

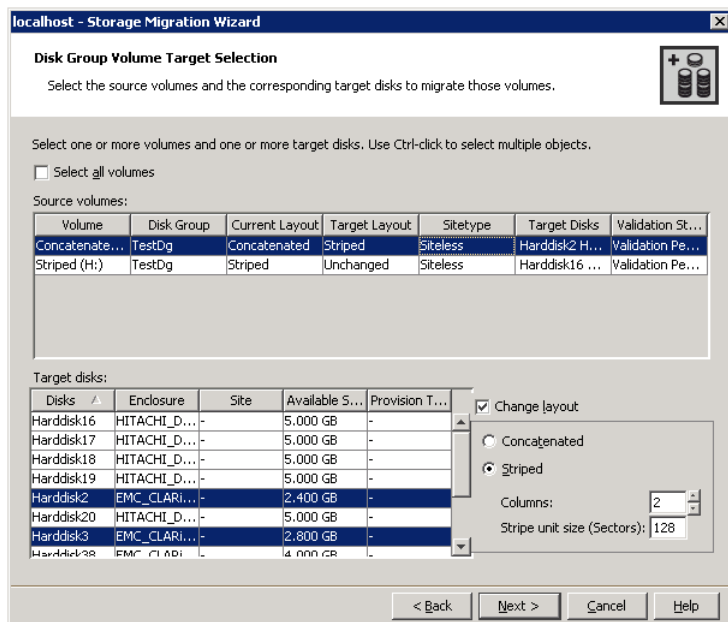
Use this procedure to migrate volumes belonging to Hyper-V virtual machines.

Ensure that the disk groups that you want to migrate are imported on the system where you run this procedure.

Administering storage migration for SFW and Hyper-V virtual machine volumes

To migrate volumes belonging to Hyper-V virtual machines

- 1 Launch VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Connect to the local system.
- 2 Launch the Storage Migration Wizard.
Right-click **localhost** and then click **Migrate Storage**.
Or, select **localhost** and then on the menu bar, click **Actions > Migrate Storage**.
- 3 Review the prerequisites on the welcome panel and then click **Next**.
- 4 On the **Storage Component Selection** panel, select **Migrate volumes used by Microsoft Hyper-V virtual machines**, and select the desired virtual machine from the drop-down list.
- 5 Click **Next** to navigate to the **Disk Group Volume Selection** panel. You can select one or more volumes for online migration. The **Select all volumes** option is selected by default. Clear the **Select all volumes** option and press Ctrl+click if you want to select one or more volumes for online migration. Click **Next**.



- 6 On the **Disk Group Volume Target Selection** panel, you can select the source volumes and the corresponding target disks to migrate volumes. The **Select**

all volumes check box is selected by default. Clear the **Select all volumes** check box to map target disks to individual volumes and **Change layout** if desired. Press Ctrl+click to select multiple target disks.

- All selected source volumes can migrate to any target disks if the **Select all volumes** option is selected.
- If you select the **Select all volumes** option, you do not have the option to change the layout. The layout can be changed only on selecting individual volumes.
- If you select the **Change layout** check box and opt to change the layout of the volume, the target layout column gets populated with the selected layout, otherwise **Unchanged** is displayed.

Click **Next** to go to the **Summary** panel.

Unless the validation status of all the volumes shows **Success**, the GUI does not navigate to the next panel. The validation status column is populated when you click **Next** on the **Disk Group Volume Target Selection** panel.

- 7 On the **Summary** panel, the mapping of the source volumes and target disks is shown based on the selection done on the previous panel. The **Storage migration summary** panel displays the name of the disk group, the source volumes, target layout and the target disks.

If the **Select all** option is selected then the summary is shown in the form of a grouping of source volumes and target disks. If the **Select all** option is unchecked, the individual selected volumes and assigned target disks are displayed on the **Summary** panel.

The **Save Summary** option on the **Summary** panel lets you save the selection summary in an HTML file.

The default summary file name is

```
StorageMigration_Summary_yyyy_MM_dd_HH_mm_ss.html.
```

The default location of the file is %allusersprofile%\veritas.

- 8 Click **Finish** to begin the volume migration.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.

- From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

Migrating data from SFW dynamic disks of one enclosure to another

Use this procedure to migrate data from SFW dynamic disks of one enclosure to another.

To migrate data from SFW dynamic disks of one enclosure to another


- 1** Launch VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
Connect to the local system.
- 2** Launch the Storage Migration Wizard.
Right-click the **Enclosures node**, and then on the menu bar click **Actions > Migrate Storage**.
Or, right-click **Enclosures Node** or an individual enclosure and then select **Migrate Storage**.
- 3** Review the prerequisites on the **Welcome** panel and then click **Next**.
- 4** On the **Storage Component Selection** panel, select **Migrate data from an array** and then from the drop-down list choose the source array enclosure.
- 5** Click **Next** to navigate to the **Enclosure Disk Selection** panel.
The disks may lie on different disk groups. Select the number of required source disks either by manually selecting one or more source disks for migration or by selecting the **Select all disks** option.

Note: Only SFW dynamic disks are available for selection as source disks.

After selecting the source disks click **Next** to continue.

- 6** On the **Enclosure Disk Target Selection** panel, select the target enclosure, as well as map source disk to the destination disk.

Administering storage migration for SFW and Hyper-V virtual machine volumes



Two options for target disk selection mode are available: **Automatic** and **Manual**.

- In the **Automatic** mode, the source, and target disks are mapped automatically.
- In the **Manual** mode, you can map source disk to target disk manually. Here there is a 1:1 mapping from source to target disk.

Click **Next** to continue after selecting the desired option.

7 The Summary panel shows the summary of the mapping.

The **Save Summary** option on the **Summary** panel lets you save the selection summary in an HTML file.

The following details are given in the summary:

- **Automatic** mode: Source Enclosure, Target Enclosure, and Source Disks: Target Disks. Here the disk grouping by automatic disk selection is shown.
- **Manual** mode: Source Enclosure, Target Enclosure, and Source Disks: Target Disks. Here the 1:1 mapping of source and target disks is shown.

The default summary file name is

StorageMigration_Summary_yyyy_MM_dd_HH_mm_ss.html.

The default location of the file is %allusersprofile%\veritas.

- 8 Review the summary and click **Finish** to begin the enclosure migration operation.
- 9 Now check the storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.

Converting your existing Hyper-V configuration to live migration supported configuration

Through the storage migration wizard, it is possible to convert your existing Hyper-V virtual machine (VM) configuration to a live migration supported configuration (one Hyper-V virtual machine per disk group).

Note: To use more than one virtual machines per disk group, either upgrade to cluster-shared disk group (CSDG) or use Cluster Volume Manager (CVM).


Before you convert or change the existing Hyper-V virtual machine configuration to a live migration supported configuration, it is presumed here that you have completed the following procedures already:

- Create disk groups and dynamic volumes.
 - See [“About creating dynamic disk groups”](#) on page 43.
 - See [“About creating dynamic volumes”](#) on page 47.
- Copy Virtual Hard Disks (VHDs) to volumes
- Create at least two Hyper-V virtual machines (VMs).

To convert a Hyper-V configuration to a live migration supported configuration

- 1** To migrate the volumes that belong to a Hyper-V virtual machine to empty disks, make sure that you have completed migrating volumes belonging to Hyper-V virtual machines.
See [“Migrating volumes belonging to Hyper-V virtual machines”](#) on page 483.
- 2** Now check storage migration status (whether successful or failed) by completing the following on the VEA:
 - Click the **Console** tab on the bottom panel of the VEA.
 - Verify that separate Subdisk move tasks are created per subdisk.
 - Click the **Tasks** tab on the VEA to check the tasks progress in the bottom panel of the console.
 - Wait for all Subdisk move tasks to complete.
 - From the **Disk View** tab verify that all selected volumes are now migrated to the selected destination disks.
- 3** After storage migration completes successfully, split the disk group into two disk groups by selecting the dynamic disk group that you created in the preceding sections. Right-click the disk group to launch the split dynamic disk group wizard
OR
On the VEA, right-click a disk group to select the split dynamic disk group option.
See [“Dynamic Disk Group Split and Join”](#) on page 272.
See [“About creating dynamic disk groups”](#) on page 43.
Click **Next**.

4 Specify the New dynamic disk group name (Say DG2).



Select **Split dynamic disk group by disks** option.

Select the disks to which Hyper-V volumes are migrated.

Click **Next**.

5 The Summary page is displayed. Click **Finish to exit the wizard.**

Now the configuration is changed to one virtual machine per disk group.

Optional Storage Foundation features for Hyper-V environments

This section includes the following topics:

- See [“About using optional Storage Foundation features in the Hyper-V parent”](#) on page 491.
- See [“Dynamic Multi-Pathing for the virtual environment”](#) on page 491.
- See [“Replicating virtual machines”](#) on page 492.
- See [“Virtual machine volume snapshots”](#) on page 493.
- See [“Campus clusters”](#) on page 494.

About using optional Storage Foundation features in the Hyper-V parent

Running Storage Foundation (SFW) in the Microsoft Hyper-V parent partition offers benefits for virtual machines (VMs) that would otherwise be unavailable at the guest level.

SFW also offers advanced features and functionality, such as multi-pathing, replication, and snapshots, as product options. More information is available on how to use the following features and on the benefits they provide when running SFW in the Hyper-V parent:

- Using Dynamic Multi-Pathing (DMP) to provide failover and load-balancing to the LUNs that host the VMs in the child partition (DMP Device Specific Modules (DSMs) option)
See [“Dynamic Multi-Pathing for the virtual environment”](#) on page 491.
- Replicating VMs between sites (Volume Replicator option)
See [“Replicating virtual machines”](#) on page 492.
- Maintaining Quick Recovery snapshots of the VMs (FlashSnap option)
See [“Virtual machine volume snapshots”](#) on page 493.

In addition, running SFW in the parent partition facilitates implementing campus clusters (stretched clusters) in the Hyper-V environment.

See [“Campus clusters”](#) on page 494.

Dynamic Multi-Pathing for the virtual environment

DMP provides an advanced multi-pathing solution for Hyper-V environments.

Multi-pathing software provides the intelligence necessary to manage multiple I/O paths between a server and a storage subsystem. This becomes even more important in virtual environments, where a single physical server hosts multiple operating system instances and applications. Loss of access to storage due to an HBA, cable, or controller failure can cause widespread impact, with potential for greater disruption than would be the case with a single physical system.

In virtual environments, a hypervisor separates the VMs from the physical hardware, making it difficult to directly map virtual host bus adapters (HBAs) in the guest to the physical HBAs that service the physical host, and to therefore have a true multi-pathing solution in the guest. Installing the required InfoScale product in the parent partition allows for true multi-pathing, providing failover, and load-balancing to the LUNs that host the VMs in the child partition.

Most multi-pathing solutions are developed by storage vendors for their specific brand of storage and, as such, generally cannot service the many types of storage

that can exist in today's heterogeneous SAN environments. In contrast, DMP is a truly heterogeneous solution. DMP fully integrates with the Microsoft Multipath I/O (MPIO) architecture. DMP includes several device-specific modules (DSMs) which provide array-specific support for a wide variety of the most popular array families available.

Replicating virtual machines


The Volume Replicator (Volume Replicator) option offered by InfoScale Storage and InfoScale Enterprise provides a means to replicate virtual machine (VM) data.

While there are host-based technologies for replicating data across distances, they're usually expensive, requiring not only more storage, but also exactly the same hardware at both ends. They can also be limited in their ability to provide a solution that accounts not only for the data, but also for the applications that access it.

Volume Replicator runs at the host level, making it possible to replicate data volumes across distances to provide a means to extended disaster recovery without requiring that the hardware be exactly the same at both ends, and generally requiring less storage. But, while this works fine to protect application data in the guest, allowing recovery at the DR site, running in the guest does nothing to protect the VM.

By installing the required InfoScale product in the parent partition, volumes that contain VHD files used for VMs and/or application data can be selectively replicated to DR sites, either synchronously or asynchronously, over an IP network. Volume Replicator uses a replicator log to store all writes to the volumes grouped together in what is known as a replicated volume group in the correct order and replicate them to the DR site, maintaining write order fidelity, and thereby, consistency. The replicated VMs remain offline at the DR site until required to be brought online, either due to an intentional migration of services from the primary site, or due to an outage at the primary site, requiring the DR site to take over services.

Figure 16-1 Volume Replicator in the parent partition



For planned outages at the primary site, perhaps for a maintenance window, the primary role can be migrated to a DR site, allowing the VMs that have been replicated to be brought online and the applications that run on them to access the data that has also been replicated to the DR site.

For unplanned outages at the primary site, operations can be moved to a DR site by a takeover operation, which turns the DR site into an active primary, allowing VMs to be brought online. Depending on the mode of replication, they can either be completely up to date or behind the previous primary. In either event, consistency is maintained and applications are able to successfully attach to their data. The primary can be migrated back to the original site when it becomes available.

Virtual machine volume snapshots

Snapshot technology is available with the Storage Foundation (SFW) FlashSnap option. Running in the guest, snapshot copies of data volumes can be created, allowing for quickly recovering from a disaster, or for off-host operations, which can occur to another virtual machine (VM) or to a physical server connected to storage that is shared with the guest when it uses pass-through disks. SFW also supports Microsoft's VSS framework for creating consistent snapshots.

However, this offers no protection against possible disaster that can occur at the VM level. If the VHD file that holds a VM is corrupted, the volume that hosts the VHD file(s) is lost, or the LUN hosting the volume that are used by the VHD file fails, snapshots in the guest are useless until the VM can be rebuilt and the application(s) reinstalled.

By running SFW in the parent partition, you have the advantage of being able to create snapshots of the volumes containing the VHDs. These snapshots can be used to quickly recover the entire VM in the event of a disaster. They can also be moved to another server and brought online to be backed up or used for other operations such as testing.

Campus clusters

As a host-based volume manager, Storage Foundation (SFW) provides the ability to mirror volumes across arrays. Clusters which rely on shared storage can be stretched beyond the confines of a single data center to a data center located at a DR site, as long as the distance between the two data centers doesn't exceed fiber channel latency limitations. These stretched clusters, also known as campus clusters, provide a level of high availability that can withstand a complete site failure.

SFW running in the parent partition can facilitate stretching of the failover cluster to another site by providing support for dynamic disks through its Volume Manager Disk Group cluster resource. With dynamic disks now available in the cluster, volumes can be mirrored across the arrays which are located in different data centers and are seen and shared by cluster nodes that are located in those respective data centers. If the active site should experience a failure, virtual machines that were running at that site can be failed over to the cluster node at the other data center, and applications running on them can be brought back online.

Microsoft Failover Clustering support

This chapter includes the following topics:

- [Overview](#)
- [Enabling Microsoft Failover Clustering after InfoScale Storage is already installed](#)
- [Configuring Microsoft Failover Clustering](#)
- [Additional considerations for SFW Microsoft Failover Clustering support](#)

Overview

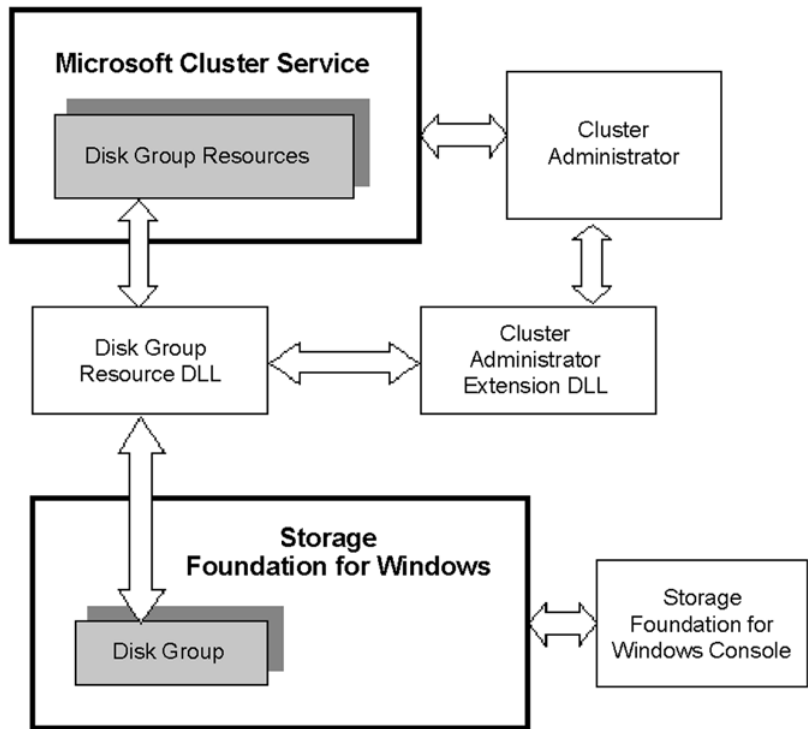
Storage Foundation supports a Microsoft Failover Clustering configuration.

Microsoft Failover Clustering provides a set of standard resources, such as physical disks, file and print shares, and network addresses and names. Each resource is accessed using a resource DLL and an extension DLL. The resource DLL provides communication between Microsoft Failover Clustering and the resource. The extension DLL allows the user to set the resource properties.

Microsoft also allows vendors to integrate product-specific resources into Microsoft Failover Clustering through the use of custom resource and extension DLLs. SFW's product-specific resource is the Storage Foundation disk group resource with its two DLLs, Disk Group resource DLL and Cluster Administrator extension DLL.

The block diagram below shows the relationship between Storage Foundation cluster components and the Microsoft Cluster Service.

Figure 17-1 Relationship between SFW and the Microsoft Cluster Service



To work in a Microsoft failover cluster environment, you must enable this feature while installing InfoScale Storage.

This installs the Storage Foundation disk group resource DLL and the Cluster Administrator extension DLL.

See [“Enabling Microsoft Failover Clustering after InfoScale Storage is already installed”](#) on page 496.

Enabling Microsoft Failover Clustering after InfoScale Storage is already installed

This section describes the steps that are needed to enable Microsoft Failover Clustering support on a server if InfoScale Storage is already installed. You need to do a "rolling install," in which you install the SFW Microsoft Failover Clustering support option on the inactive node or nodes of the cluster first. Then you make

Enabling Microsoft Failover Clustering after InfoScale Storage is already installed

the active cluster node inactive by moving the cluster resources to another node and install on that node.

Note: After completing the steps that are provided below, you must ensure that the Volume Manager Disk Group (VMDg) resource has been registered. If not, then you need to register it manually. For information about registering the resource manually, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

To enable Microsoft Failover Clustering option when SFW is installed

- 1 First, you must have a Microsoft failover cluster setup and have failover clusters running on each server where you want to install the SFW Microsoft Failover Clustering option. SFW also must be running on each server in the cluster.
- 2 On the inactive node of the cluster, select **Add or Remove Programs** in the Windows Control Panel.
- 3 From the menu that is located on the left side of the Add or Remove Programs window, select **Change or Remove Programs**.
- 4 Highlight the SFW Server Components entry and click **Change**.
- 5 The installer window appears. Select **Add or Remove** to add or remove features. Click **Next** to continue.
- 6 The Option Selection window appears. Check the **Cluster Option for Microsoft Failover Cluster** check box.

If the option is grayed out, add its license key with the following:

- Click the **Add License** link that is located at the far right of the window to add a license key for an option.
- In the pop-up window that appears, enter the license key for the option and click **OK**. Click the check box to add the option.

Click **Next** to continue.

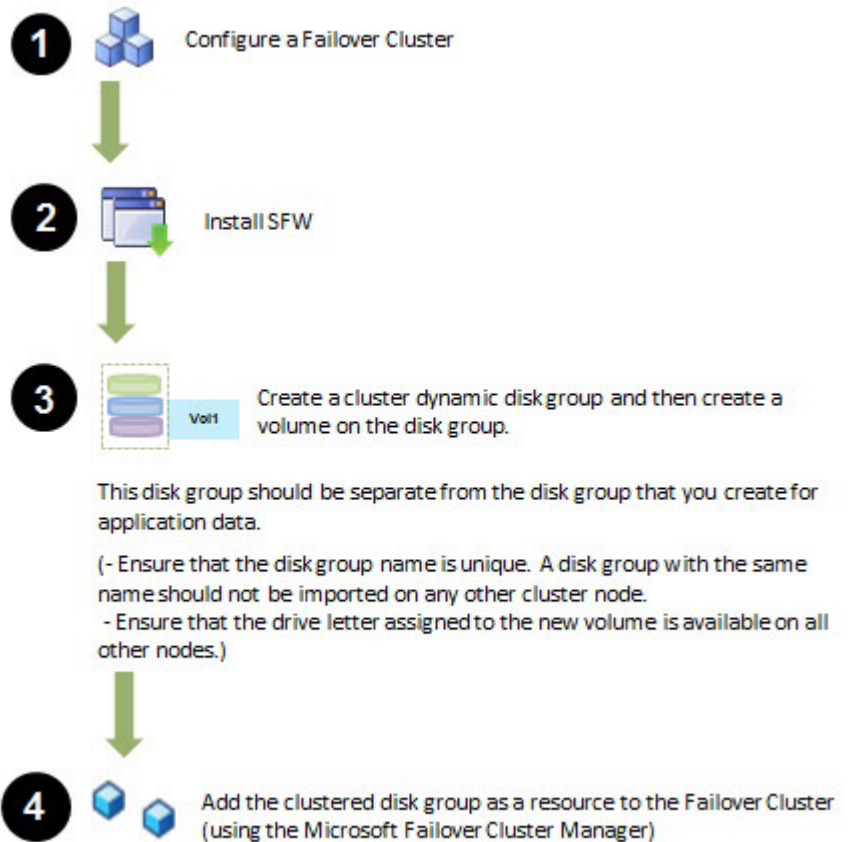
- 7 The **Validation** window appears. The installer checks for prerequisites for the selected systems and displays the results. Review the information and click **Next** to continue.
- 8 The **Summary** window appears. Review the information and click **Update** to begin the product update.

The **Update Status** window appears. Status messages and the progress of the update are displayed. At the completion of the update, click **Next** to continue.
- 9 A report summarizing the update appears. Review the information and click **Next** to continue.

- 10** The **Thank You** window appears when the feature has been added or removed. Click **Finish** to complete the process.
- 11** A message box appears. Click **Yes** to restart your system and complete the update.
- 12** Complete the installation by repeating the sequence on any remaining nodes.
For the active node, use the **Move** command with Windows Server's Failover Cluster Manager snap-in to move the cluster resources to another node before installing.

Configuring Microsoft Failover Clustering

To configure Microsoft Failover Clustering, you must perform the following steps:



Task

Step 1

Step 2

Step3

Reference

For details about configuring Failover Cluster, refer to Microsoft documentation.

For details about installing the product, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

See [“Create a cluster dynamic disk group”](#) on page 500.

See [“Creating a dynamic volume”](#) on page 48.

Task

Step 4

Reference

See [“Adding the clustered disk group as a resource to the Failover Cluster”](#) on page 500.

Note: SFW normally does not let you create a cluster dynamic disk group with the disks that are on the same bus as the cluster node’s system or boot disk. To facilitate booting from a SAN, you can remove this limitation by running a `vxclus` command from the SFW command line interface.

Create a cluster dynamic disk group

To create a cluster disk group, See [“About creating dynamic disk groups”](#) on page 43.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Please note the following consideration in this process:

- Setting up the cluster disk group on a Windows Server operating system creates physical disk resources for all the basic disks on the shared bus. Later, you create resources for the SFW cluster disk groups. Before doing so, you must remove any physical disk group resources for the disks that are used in the cluster disk groups. Otherwise, a reservation conflict occurs.

Adding the clustered disk group as a resource to the Failover Cluster

The following steps describe how to add a Storage Foundation clustered disk group as a resource to the Failover Cluster.

To add the clustered disk group as a resource to the Failover Cluster

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

To launch the Failover Cluster Manager snap-in, click **Failover Cluster Manager**.

- 2 Expand the cluster from the tree view in the left pane to display **Roles**.
- 3 Right-click **Roles**, and then select **Create Empty Role**.

The new role is created under the **Roles** pane in the center.

- 4 Right-click the role that you created, and then select **Properties**.
 The **New Role Properties** window appears.
- 5 Enter a name for the new role in the **Name** field.
- 6 Select the desired cluster nodes and priority, and then click **OK** to continue.
- 7 To add a Volume Manager Disk Group (VMDg) resource, right-click the role, and then select **Add Resource > More Resources > Volume Manager Disk Group**.

In the pane below the **Roles** pane, click the **Resources** tab to see information about the new VMDg resource.

- 8 In the pane below the **Roles** pane, click the **Resources** tab, right-click **New Volume Manager Disk Group**, and then select **Properties**.
 The **New Volume Manager Disk Group Properties** window appears.
- 9 On the **General** tab, enter a name for this new resource in the **Name** field.
- 10 On the **Properties** tab, view and modify the private properties of this resource, and then click **OK** to continue.

The table in the properties panel gives the following details:

Name	Value
DiskRunChkDsk	Specify whether SFW should perform a file system check on the disks each time the cluster resource is brought online. If enabled, SFW automatically cleans the file system by running the <code>chkdsk /x /f</code> command on the volume that is brought online, if dirty bit is set on that volume. The value 1 (enabled) indicates that SFW performs the file system check and the value 0 (disabled) indicates that it does not. The default is 0 (disabled).
FastFailover	Specify whether the VMDg resource should be enabled for fast failover. Specify "true" for enabling it and "false" for disabling it. The default value is "true".
DiskGroupName	Specify the name of the SFW cluster disk group that you want to configure as a cluster resource.

Note: A VMDg resource is not marked as online until all volumes in it are verified using `chkdsk` utility, if the value is set to 1.

- 11 The **Summary** tab displays the new cluster resource. The cluster resource's status should be online. If it's not, then right-click it, and then select **Bring Online** to bring the resource online.
- 12 In the tree-view of Failover Cluster Manager, expand the **Storage** node to review the attributes of the new cluster resource.

Notes:

- SFW uses the Windows temporary folder (typically `C:\Windows\temp`) to store the output of the `chkdsk` command.

If there is not enough disk space or if the logged-on user does not have write access to the temporary directory, the Volume Manager Disk Group (VMDg) resource may fail due to a `chkdsk` error.

The cluster log shows the following message:

```
ERR [RES] Volume Manager Disk Group <FORCHKDSK>: RunChkDsk: Failed to get the temporary file.
```

You need to run `chkdsk` manually to bring the resource online.

- `chkdsk` fails to execute on a disk group under replication and does not come online.
If `chkdsk` is run on a volume that is under replication, cluster logs display the following error message:
Windows cannot run disk checking on this volume because it is write protected. In such a scenario, replication locks the volume, as a result `chkdsk` cannot be executed, and the VMDg resource is marked as failed because `chkdsk` has failed.
Therefore, the VMDg does not come online.
Do not enable `chkdsk` on a disk group that is under replication.
Remove a disk group out of replication and run `chkdsk` manually.
- If the VMDg resource has a dirty volume, then Windows shows a pop-up while the volume mounts. Set the `chkdsk` property of the VMDg resource under Microsoft Failover Clustering to '1' to perform automatic verification of that volume.

Displaying the cluster disk group resource properties

The following displays the cluster disk group resource properties.

To display cluster disk group resource properties

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

To launch the Failover Cluster Manager snap-in, click **Failover Cluster Manager**.

- 2 Expand the tree view in the left pane to display the **Storage** node.

Expand the **Storage** node, and then select **Disks**. The cluster disk group resource is displayed in the **Disks** pane.

The Summary of Storage window appears displaying the cluster disk group resource.

- 3 Right-click the cluster disk group resource and select **Properties** from the context menu.

The **Properties** window appears where you can view the properties of the cluster disk group resource.

Additional considerations for SFW Microsoft Failover Clustering support

This section contains additional information that is important in working with Microsoft Failover Clustering and Storage Foundation for Windows.

Note the following considerations:

- When a cluster disk group resource is offline or a cluster disk group that is not a failover cluster resource is in a Deported state, it is not protected from access by other computers. For maximum data protection, keep Volume Manager Disk Group resources online. Note that the SFW disk group resources still retain the "Volume Manager" name.
- When using the Windows Server's Failover Cluster Manager snap-in to create a disk group resource, the Volume Manager Disk Group Parameters screen might not list all the available Storage Foundation for Windows cluster disk groups in the drop-down list. If this happens, exit the New Resource wizard and use the Windows Server's Failover Cluster Manager snap-in to select the cluster group to which the resource is to be assigned. Next, move the cluster group to the cluster node where the Storage Foundation for Windows cluster disk group is currently online. Then create the Storage Foundation for Windows disk group resource.
- Under the following circumstances, the VEA Disk View may not reflect the latest state of the disk(s) until a refresh is performed:

- When you change the state of a cluster disk resource on one node and try to view the disks under this resource from another node on the same cluster.
- When you change the state of a cluster disk resource on one node and try to view the disks under this resource from a remote computer.
- SFW support of the Microsoft Failover Clustering environment allows the selection of SCSI-2 reservation mode or SCSI-3 reservation mode. Selecting the type of SCSI support for the Microsoft Failover Clustering environment is done by using the System Settings portion of the SFW Control Panel. When selecting the type of SCSI support in a Microsoft Failover Clustering environment, it is important to know if your storage arrays support SCSI-3. SFW SCSI-3 clustering support does not let you mix the storage arrays that support SCSI-3 with the storage arrays that cannot. In a situation of mixed storage arrays, you must use SFW SCSI-2 clustering support. Refer to the HCL for arrays that support SCSI-3.

Note: Veritas maintains a hardware compatibility list (HCL) for Storage Foundation and High Availability Solutions for Windows products on the Veritas support Web site. Check the HCL for details about your storage arrays before selecting the type of SCSI support in a Microsoft Failover Clustering environment.

After selecting the type of SCSI support, you must issue the following CLI commands to complete the setting on your system:

- `net stop vxsvc`
- `net start vxsvc`

Note: If a cluster disk group is imported on the system, you must deport or move the cluster disk group to another system before issuing these CLI commands.

If SFW SCSI-2 clustering support is selected and Active/Active load balancing is desired, the SCSI-3 Persistent Group Reservations (SCSI-3 PGR) support mode must be enabled for the DMP DSM.

- A cluster dynamic disk group that is part of the cluster resources cannot be a source disk group for a join command. However, it can be a target disk group for the command.
- **Change in Bringing a Two-Disk Cluster Group Online**
 In earlier versions of Volume Manager for Windows, it was possible to bring a two-disk cluster disk group online when only one disk was available. If a cluster were to lose all network communication, this allowed the disk group to be brought online on two cluster nodes simultaneously, with each node owning a single

disk, possibly resulting in data loss or a partitioned cluster. Though the likelihood of this situation occurring is slim for most customers, the consequences if it does happen may be severe. However, this is no longer supported with recent versions of Volume Manager and it is not possible to bring a two-disk cluster disk group online in Volume Manager unless it complies with the normal majority algorithm which means both disks must be available.

The normal majority algorithm is $(n/2 + 1)$.

- You are not allowed to deport a cluster disk group that is also a Volume Manager disk group resource for Microsoft Failover Clustering.

- Connecting to a Cluster Node

If you connect to a computer from the VEA GUI using the virtual name or the virtual IP address, the VEA GUI displays the computer name of the cluster node that currently owns the virtual name and IP resources. Therefore, it is not recommended to use the virtual name or virtual IP address when connecting and administering a cluster node through SFW HA.

Instead, use the actual computer name or the IP address of the cluster node.

- Dynamic Multi-Pathing (DMP) does not support using a basic disk as a cluster resource under Microsoft Failover Clustering.

Failover may not function properly when using Dynamic Multi-Pathing with a Microsoft Failover Clustering basic disk cluster resource. Refer to Tech Note 251662 on the Veritas Support site for details.

If you want to use Dynamic Multi-Pathing with SFW and Microsoft Failover Clustering, you must convert any Microsoft Failover Clustering basic disk cluster resources to dynamic disk cluster resources before activating Dynamic Multi-Pathing. The initial setup of Microsoft Failover Clustering requires that you use a basic disk as the quorum disk. Once InfoScale Storage is installed, you should upgrade the basic disk to dynamic by including it in a dynamic cluster disk group and then convert the quorum resource from a basic disk resource to a dynamic disk resource.

Note: DMP DSMs do not support an Active/Active setting in a Microsoft Failover Clustering environment when a quorum disk is a basic disk.

- Cluster dynamic disk groups that contain iSCSI disks are not set up for persistent login on all nodes in the cluster.

SFW ensures that the iSCSI targets of cluster dynamic disk groups that contain iSCSI disks are configured for persistent login. If the persistent login is not configured for the target, SFW automatically configures it.

Cluster dynamic disk groups that contain iSCSI disks are only automatically configured for persistent login on the node where they were created. The other

nodes in the cluster are not enabled for persistent login. You need to manually set up the persistent login for each of the other nodes in the cluster.

- Copying the policy file, VxVolPolicies.xml, to Another Node
If the second node is configured the same as the first and if the first node's policy settings for Automatic Volume Growth are to be maintained on the second node, you need to copy the VxVolPolicies.xml file of the first node to the second node. Copy the VxVolPolicies.xml file to the same path location on the second node as its location on the first node. The default path of the VxVolPolicies.xml file is Documents and Settings\All Users\Application Data\Veritas.
More information about the policy file is available.
See "[Configuring Automatic volume growth](#)" on page 197.
- More information about using SFW and Microsoft Failover Clustering in a shared cluster environment with the FlashSnap off-host backup procedure is available. See "[Using Dynamic Disk Group Split and Join with a cluster on shared storage](#)" on page 282.
- If you install the Microsoft Failover Clustering feature on a server on which InfoScale Storage for Windows is already installed, then you must manually restart Veritas Enterprise Administrator Service (VxSvc) by running the following commands:
 - `net stop vxsvc`
 - `net start vxsvc`

Configuring a quorum in a Microsoft Failover Cluster

This chapter includes the following topics:

- [Overview](#)
- [SFW support for dynamic quorum resource](#)
- [Configuring Failover Cluster quorum using SFW](#)
- [Microsoft Failover Cluster quorum arbitration time settings](#)

Overview

Microsoft Failover Cluster provides an option to configure a quorum for the cluster. During the cluster configuration you are required to assign a separate disk for configuring the quorum. The disk that is assigned for quorum stores a copy of cluster configuration data.

A cluster uses the quorum for the following tasks:

- Define which cluster node should be active at any given time.
- Define on which cluster node the service group should be brought online, in the event of a split-brain.

By default, the failover cluster uses a single physical disk resource and a basic disk volume resource for a quorum.

If the physical disk resource fails or becomes corrupted, the cluster becomes unusable. To avoid this failure you must use a dynamic disk group resource. Unlike a physical disk resource which contains a single disk, a dynamic disk group resource contains more than one disk. A dynamic disk group resource provides a high level of redundancy by allowing the mirroring of disks.

SFW support for dynamic quorum resource

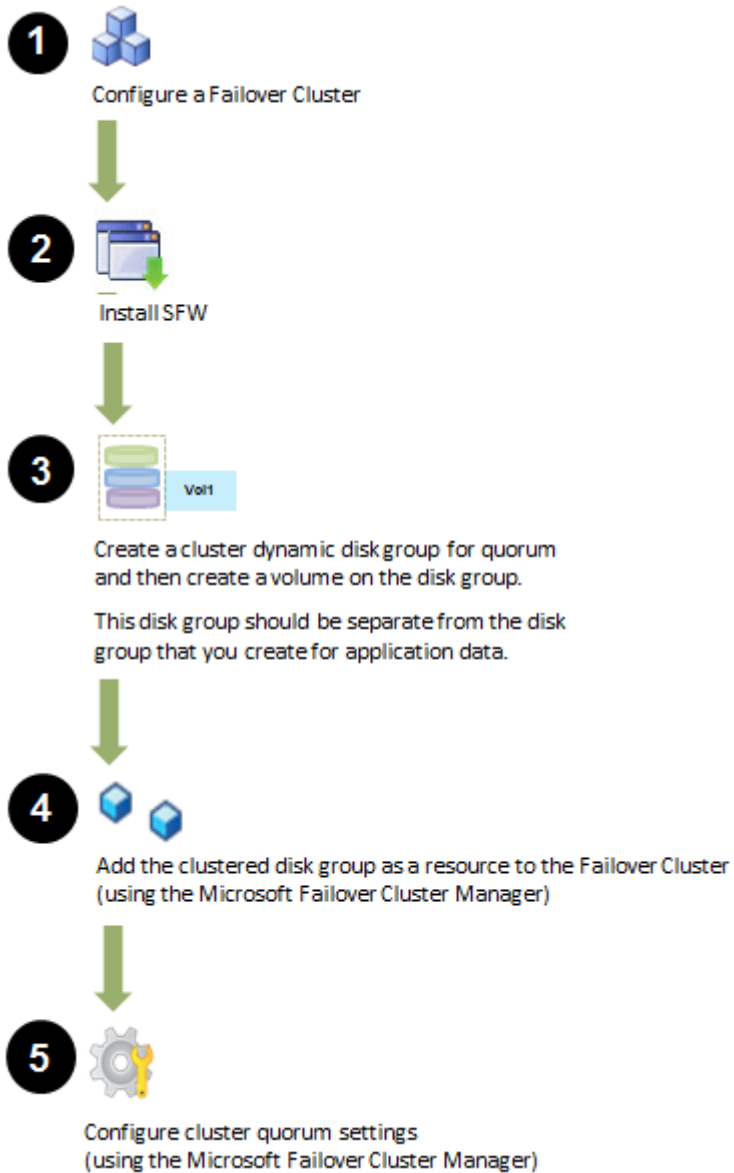
Storage Foundation (SFW) allows Failover Cluster to use dynamic disks as clustered dynamic disk group resources. This lets you make your Failover Cluster quorum highly available using host-based mirroring.

Veritas recommends the following:

- Use at least three disks for a clustered dynamic disk group resource because a cluster disk resource cannot be brought online unless a majority of disks are available.
- In case of Cluster Volume Manager (CVM) clusters, use physical disks for configuring a quorum.

Configuring Failover Cluster quorum using SFW

The following figure depicts how to configure Failover Cluster quorum.



Task

Step 1

Reference

For details about configuring Failover Cluster, refer to Microsoft documentation.

Task	Reference
Step 2	For details about installing the product, refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i> .
Step 3	See “Create a cluster dynamic disk group” on page 500. See “Creating a dynamic volume” on page 48.
Step 4	See “Adding the clustered disk group as a resource to the Failover Cluster” on page 500.
Step 5	See “Configuring cluster quorum settings” on page 510.

Configuring cluster quorum settings

The following procedure outlines the steps for configuring a clustered dynamic disk group resource.

Note: You must perform these tasks from a cluster node where all the cluster resources are online. If the clustered dynamic disk group resource is online on a node other than the one where rest of the cluster resources are online, you must move the clustered dynamic disk group resource to the other node and then perform these tasks.

To move the resource, right-click the resource and from the context menu select, **Move > Best Possible Node**.

To configure a clustered dynamic disk group resource

- 1 From the **Start** menu (the **Start** screen on Windows 2012 operating systems), click **Administrative Tools**.

To launch the Failover Cluster Manager snap-in, click **Failover Cluster Manager**.

- 2 Right-click the cluster node and in the context menu, select **More Actions > Configure Cluster Quorum Settings**.

The **Configure Cluster Quorum** Wizard appears.

Click **Next** to continue.

- 3** On the **Select Quorum Configuration Option** panel, select **Select the quorum witness**.
Click **Next** to continue.
- 4** On the **Select Quorum Witness** panel, select **Configure a disk witness** and then click **Next**.
- 5** On the **Configure Storage Witness** panel, select the disk group that is added for cluster quorum, and then click **Next**.
- 6** Review the report that the wizard generates and click **Finish** to complete the operation.

Confirm that the VMDg resource moves to “Cluster Group” along with Network Name and IP resources.

The wizard configures the cluster quorum settings.

The Volume Manager Disk Group that is created for cluster quorum is now added in the Failover Cluster as a quorum resource.

The quorum data is mirrored on all the disks that are part of the Volume Manager Disk Group. Even if any of the disks fail, the quorum data is available on the remaining disks.

Microsoft Failover Cluster quorum arbitration time settings

When you select the SFW Microsoft Failover Clustering support option during installation, a message appears to remind you that the maximum quorum arbitration time can be adjusted to ensure optimal functioning of SFW dynamic volumes with Microsoft Failover Clustering.

The quorum arbitration time setting is used to set the limit of the time period that is allowed for quorum arbitration. Quorum arbitration is the process that occurs when the controlling node of the cluster is no longer active and other nodes of the cluster attempt to gain control of the quorum resource and thus control of the cluster.

The advantage of a dynamic quorum is that it can be mirrored to provide fault tolerance for the quorum resource.

Note: In Windows Server 2012, to use the Failover Cluster Command Interface (cluster.exe) feature for Microsoft Failover Clustering, you must select the **Failover Cluster Command Interface** option under the **Failover Clustering Tools** feature during Windows Server 2012 installation. For more information, refer to the Microsoft documentation.

The optimal quorum arbitration time setting for a dynamic quorum is:

Maximum time - 120 seconds

To specify this setting, type the following command in the command line:

```
cluster.exe /prop QuorumArbitrationTimeMax=120
```

If you do not plan to use a dynamic quorum and instead want to retain the quorum resource on a basic disk, then you may want to specify the default settings that Windows server recommends for a basic quorum.

The Windows server quorum arbitration time setting for a basic quorum is:

Maximum time - 60 seconds

To reset the setting, type the following command in the command line:

```
cluster.exe /prop QuorumArbitrationTimeMax=60
```

The advantage of changing the setting back is that it speeds up the quorum arbitration process.

Implementing disaster recovery with Volume Replicator

This chapter includes the following topics:

- [Overview](#)
- [Summary of the steps for setting up a disaster recovery solution with Volume Replicator](#)
- [Volume Replicator terms](#)

Overview

Volume Replicator (Volume Replicator) option is a data replication tool designed to maintain a consistent copy of application data at a remote site as part of an effective disaster recovery plan. In the event that the data center is down, the application data is immediately available at the remote site, and the application can be restarted at the remote site.

Volume Replicator works as a fully integrated component of InfoScale Storage or InfoScale Enterprise. Volume Replicator benefits from the robustness, ease of use, and high performance of InfoScale Storage or InfoScale Enterprise and, at the same time, adds replication capability. Volume Replicator can use existing InfoScale Storage or InfoScale Enterprise configurations, with some restrictions. Any application, even with existing data, can be configured to use SFW transparently.

For detailed information about Volume Replicator, refer to the *Volume Replicator Administrator's Guide*.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support Volume Replicator.

Volume Replicator feature highlights

Volume Replicator has many feature highlights.

They include the following:

- Supports replication of data over any IP network, LAN, or WAN.
- Runs on all storage hardware that Storage Foundation supports.
- Supports replication over a firewall.
- Provides volume-level replication of application or file system data, including support of commercial database management systems.
- Replicates the data in asynchronous or synchronous mode, ensuring complete data integrity and consistency in either mode.

How Volume Replicator works

Volume Replicator's purpose is to replicate data from a primary site to one or more secondary sites. It does this by using a replicated volume group (RVG) within an SFW disk group as the unit of replication.

The following is a summary of how Volume Replicator works:

- Through the Volume Replicator software, the volumes to be replicated on the primary site are identified as part of an RVG, which consists of one or more volumes in a SFW disk group. If you have multiple disk groups with volumes to be replicated, each disk group must have a separate RVG. It is possible to have more than one RVG per disk group.
- With each RVG, a Replicator Log volume is also set up. The Replicator Log volume at the primary site holds the writes that are to be sent to the secondary site.
- A corresponding RVG and Replicator Log volume at the secondary site are also set up.

An identical disk group and volume setup is created on the secondary site. The disk groups and volumes must be of the same size and have the same names as those on the primary site. The volumes do not have to be the same volume type.

The Replicator Log volume on the secondary site must have the same name as on the primary site, but its size can differ. However, Veritas recommends that the two log volumes be the same size.

The secondary site Replicator Log is held in reserve so that it can be used if the primary site goes down or has to be migrated and the secondary site needs to become the new primary site.

- The RVG at the primary site and the corresponding RVG at the secondary site are called a Replicated Data Set (RDS). Most Volume Replicator commands operate on an RDS. Normally, you can perform Volume Replicator operations from any host in an RDS.
- Once the Volume Replicator components are properly installed and configured, replication starts.

Volume Replicator uses the Replicator Log volume on the primary site to track all the writes to the application or file system in the order that they were received and then transmits the writes to the secondary site. Each write to a data volume under an RVG on the primary site generates two writes: The first one is sent to the Replicator Log, and when that is complete, the other is sent to the application data volumes and to the secondary site at the same time.

When the secondary system receives a write, it sends an initial acknowledgment of the receipt back to the primary site, even before the write is committed to disk. This is called the "Network Acknowledgment." Once the secondary commits the write to disk, a second acknowledgment, called the "Data Acknowledgment," is sent to the primary system. The Replicator Log volume on the primary system discards the write when it receives the Data Acknowledgment.

Replication is a unidirectional process. The updates on the primary host are sent to the secondary host, but access to the data at the secondary host or hosts is read-only on the replication volumes.

- The three modes of replication — synchronous, asynchronous, and synchronous override — work as follows:
 - The synchronous mode waits until the Network Acknowledgment has been received from the secondary host before it completes the write to the application. Thus, the primary and the secondary have the same data.
 - The asynchronous mode completes the application write after it has been written to the primary Replicator Log volume.
If the primary site goes down, there may still be some writes that were not yet received at the secondary site. This mode has better performance but with a risk of some data loss.
 - The synchronous override is a mode of replication that is synchronous as long as the network is available, but when the network becomes unavailable, the replication is continued in the asynchronous mode.
- If a disaster occurs on the primary site and its data is destroyed, a secondary host can take over the role of the primary host to make the data accessible. You can then restart the application on that host.

- You can also manually migrate the role of a healthy primary host to a secondary host when the application that is involved in replication is inactive. You may want to do this for maintenance purposes.

Summary of the steps for setting up a disaster recovery solution with Volume Replicator

This section provides a high-level summary of the steps for setting up a Volume Replicator disaster recovery solution with SFW. For more detailed information, refer to the *Volume Replicator Administrator's Guide*.

These instructions do not give all the steps for setting up a cluster with Volume Replicator. For full details on clustering steps with SFW HA, see the *Storage Foundation and High Availability Solutions Solutions Guide*. You can refer to application-specific Solutions guide for more details on how they work with SFW HA. For example, if you use SQL Server with SFW HA, refer to the *Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.

All the guides are included on the product disc.

To set up a cluster with Volume Replicator:

- 1 Set up and configure the hardware at each site.

For Volume Replicator, installation requires at least two servers running SFW with a network connection between them: one as the primary at the main site and the other as the secondary at a second site. Optionally, you can have additional secondary sites.

If you plan to use clustering with SFW HA or SFW and Microsoft Failover Clustering, you have several configuration choices. Probably the most common one is to have four servers, two each clustered together at each site. You can also have a cluster with two servers, one node at the primary site and the other at the secondary. Another variation is for the primary site to have a two-node cluster while the secondary site has a standalone server for replication.

- 2 Install the operating system and configure the network.
- 3 Install InfoScale Storage or InfoScale Enterprise.

InfoScale Enterprise includes Global Cluster Option (GCO), that facilitates replication support after a site failure.

Summary of the steps for setting up a disaster recovery solution with Volume Replicator

- 4 After the installation completes, the Volume Replicator Security Service Configuration Wizard (VxSAS) is launched.

Follow the instructions in the Volume Replicator documentation to configure this service, which is required for Volume Replicator commands to work properly.

- 5 Create one or more disk groups in SFW or SFW HA on your primary system. Then create the volumes within each disk group.

See [“About creating dynamic disk groups”](#) on page 43.

See [“About creating dynamic volumes”](#) on page 47.

You are creating the volumes to be replicated. You can create the volume for the Volume Replicator Replicator Log now or wait until you run the Volume Replicator wizard for setting up RDS on the system.

The wizard lets you either identify an existing volume for the log volume or to have the wizard bring up the screen for creating the volume so you can create the volume at that time.

- 6 If you plan to use an application on your system, you need to install and configure the application at this point.
- 7 On your secondary system, create a duplicate of the disk groups and volumes that are on the primary system. You can do this step in one of the two following ways:
 - Repeat the same manual procedures to set up the disk groups and volumes that you used previously on the primary system.
 - Wait until you run the Volume Replicator RDS Wizard and allow Volume Replicator to make a duplicate of the disk groups and volumes for you.

The secondary disk groups and volumes should have the same names as the primary disk groups and volumes. The data volumes on the secondary site should be the same size as the corresponding data volumes on the primary site. The log volume on the secondary can be a different size, but Veritas recommends that the sizes be the same.

If you have installed an application on the primary site, it also needs to be installed on the secondary site.

- 8 If you have VCS or Microsoft Failover Clustering clusters, create virtual IP resources for replication.

You need to create a virtual IP resource that is used for replication for each replicated disk group. In addition, if you are using VCS, you need to specify a NIC resource. This NIC resource is not needed for a Microsoft failover cluster. Create the IP resource on both the primary and secondary sites. Each site needs its own IP address. This virtual IP should be linked to the NIC that is used for the server at each site.

- 9 Set up the Replicated data sets for Volume Replicator.

Refer to the *Volume Replicator Administrator's Guide* or the Storage Foundation Solutions guides for detailed instructions.

Volume Replicator terms

This section provides the definitions of the most commonly used Volume Replicator terms for reference purposes.

Replicated Volume Group (RVG)

An RVG is made up of one or more volumes in an SFW disk group. The updates that are made on the RVG on the primary host are sent to a configured secondary host. Thus, there is a corresponding RVG with a disk group of the same name and volumes with the same names. The data volumes should be the same size, but Replicator Log volume sizes can differ. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG on the primary host.

An RVG within a disk group is the container for replication, so if you have multiple disk groups, you need to create a separate RVG for each disk group. It is possible to have more than one RVG in a disk group; however, the RVG cannot span across disk groups.

Replicated Data Set (RDS)

An RVG on the primary host and the corresponding duplicate RVG on the secondary host or hosts make up a Replicated Data Set (RDS).

Replicator Log

Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any data writes that are sent to the secondary site. The Replicator Log on the secondary site is held in reserve so that it can be

used if the primary site becomes nonfunctional and the secondary site takes over the role of primary site. The logs at the two sites must have the same name; however, the sizes of the logs can differ. Veritas recommends having Replicator Log volumes of the same size at the primary site and the secondary site.

Replication Modes

The three modes of replication —synchronous, asynchronous, and synchronous override — work as follows:

- The synchronous mode waits until the Network Acknowledgment has been received from the secondary host before it completes the write to the application. Thus, the primary and the secondary have the same data.
- The asynchronous mode completes the application write after it has been written to the primary Replicator Log volume.
If the primary site goes down, there may still be some writes that were not yet received at the secondary site. This mode has better performance but with a risk of some data loss.
- The synchronous override is a mode of replication that is synchronous as long as the network is available, but when the network becomes unavailable, the replication is continued in the asynchronous mode.

Troubleshooting and recovery

This chapter includes the following topics:

- [Overview](#)
- [Using disk and volume status information](#)
- [SFW error symbols](#)
- [Resolving common problem situations](#)
- [Commands or procedures used in troubleshooting and recovery](#)
- [Additional troubleshooting issues](#)
- [CVM issues](#)
- [Tuning the VDS software provider logging](#)

Overview

This chapter describes ways to use Storage Foundation (SFW) to troubleshoot and recover from problems.

Note: It is recommended that you use a Volume Manager Disk Group (VMDg) resource for creating Microsoft file shares.

Using disk and volume status information

If a disk or volume fails, it is important to repair the disk or volume as quickly as possible to avoid data loss. Because time is critical, Storage Foundation makes it easy for you to locate problems quickly. In the **Status** column of the **Disks** tab view or **Volumes** tab view, you can view the status of a disk or volume. You can also see indications of abnormal status in the tree view or the **Disk View** tab. If the status is not Healthy for volumes, Imported for dynamic disks, or Online for basic disks, use this section to determine the problem and then fix it.

See [“Disk status descriptions”](#) on page 521.

See [“Volume status descriptions”](#) on page 524.

Note: When there is a problem with a disk or volume, you can also look at the **Events** tab view for messages about the problems that have occurred.

Disk status descriptions

One of the following disk status descriptions always appears in the **Status** column of the disk in the right pane of the console window. If there is a problem with a disk, you can use this troubleshooting chart to diagnose and correct the problem. Caution: The suggested actions may bring the disks to an Imported status, but they do not guarantee data integrity.

Table 20-1 Disk Status Descriptions

Status	Meaning	Action Required
Imported	The disk is accessible and has no known problems. This is the normal disk status for dynamic disks.	No user action is required.
Online	The disk is accessible and has no known problems. This is the normal disk status for basic disks.	No user action is required.
No Media	No media has been inserted into the CD-ROM or removable drive. Only CD-ROM or other removable disk types display the No Media status.	Insert the appropriate media into the CD-ROM or other removable drive, and the disk status becomes Online. If you do not see the disk status change immediately, use the Refresh command to refresh the GUI.

Table 20-1 Disk Status Descriptions (*continued*)

Status	Meaning	Action Required
Foreign	Only dynamic disks display this status. There are three situations where a disk may be marked as Foreign:	See “Bring a foreign disk back to an online state” on page 531.
	<p>Situation 1</p> <p>The disk was created as a dynamic disk on another computer and has been moved to your computer and has not been set up for use.</p>	<p>Situation 1</p> <p>Use the Import Dynamic Disk Group dialog box to make the disk group available for use.</p> <p>In the Import Dynamic Disk Group dialog box, be sure to click the check box to clear the host ID of the other system.</p>
	<p>Situation 2</p> <p>The disk contains a secondary disk group (that is, a disk group other than disk group containing the computer’s boot or system disk) and you have a dual-boot system. When you switch between operating systems, the disk with a secondary disk group is marked as Foreign and does not auto-import.</p>	<p>Situation 2</p> <p>Use the Import Dynamic Disk Group dialog box to make the secondary disk group available for use. The primary disk group is imported automatically when you switch between operating systems.</p> <p>In the Import Dynamic Disk Group dialog box, be sure to click the check box to clear the host ID of the other system.</p>
	<p>Situation 3</p> <p>The disk was originally created on this computer. It was moved or deleted. Now you want to bring it back to this computer as a member of the disk group in which it was originally created.</p>	<p>Situation 3</p> <p>Use the Merge Foreign Disk command to restore the disk as a member of its former disk group.</p>
No Disk Signature	Displays for new disks. The disk cannot be used because it has no signature.	Right-click on the disk and select Write Signature from the menu. The disk type changes to Basic Disk and the disk can be accessed or upgraded.

Table 20-1 Disk Status Descriptions (*continued*)

Status	Meaning	Action Required
Offline	Only dynamic disks display this status. There are two situations where a disk may show a status of Offline:	
	<p>Situation 1</p> <p>The disk was part of the system's disk configuration but currently cannot be found.</p>	<p>Situation 1</p> <p>Make sure that the disk is connected to the computer. Then use Rescan to bring the disk online.</p> <p>See "Bring an offline dynamic disk back to an imported state" on page 529.</p>
	<p>Situation 2</p> <p>The disk is not accessible. The disk may be corrupted or intermittently unavailable. An error icon appears on the offline disk.</p> <p>If the disk status is Offline and the disk's name changes to Missing Disk (#), the disk was recently available on the system but can no longer be located or identified.</p>	<p>Situation 2</p> <p>Make sure that the disk is connected to the computer. Then use Rescan to bring the disk online.</p> <p>See "Bring an offline dynamic disk back to an imported state" on page 529.</p>
Disconnected	Displays for a dynamic disk when the system can no longer find the disk. The name of the disk becomes "Missing Disk."	Reconnect the disk.
Import Failed	The import of the dynamic disk group containing the disk failed. All disks in a dynamic disk group that failed to import show this status.	Examine your configuration to determine what the problem might be.

Table 20-1 Disk Status Descriptions (*continued*)

Status	Meaning	Action Required
Failing	Failing is a secondary message shown in parentheses after a disk status. This status indicates that I/O errors have been detected on a region of the disk. All the volumes on the disk display Failed, Degraded, or Failing status, and you may not be able to create new volumes on the disk. Only dynamic disks display this status.	Right-click on the failing disk and select Reactivate Disk to bring the disk to Online status and all its volumes to Healthy status.

Volume status descriptions

One of the following volume status descriptions always appears in the graphical view of the volume and in the Status column of the volume in list view. If there is a problem with a volume, you can use this troubleshooting chart to diagnose and correct the problem.

Warning: The suggested actions may bring the volumes to a Healthy status, but they do not guarantee data integrity.

Table 20-2 Volume Status Descriptions

Status	Meaning	Action Required
Healthy	The volume is accessible and has no known problems. This is the normal volume status. Both dynamic volumes and basic volumes display the Healthy status.	No user action is required.

Table 20-2 Volume Status Descriptions (*continued*)

Status	Meaning	Action Required
Resynching	<p>The volume's mirrors are resynchronized so that both mirrors contain identical data. Both dynamic and basic mirrored volumes display the Resynching status.</p> <p>When resynchronization is complete, the mirrored volume's status returns to Healthy. Resynchronization may take some time, depending on the size of the mirrored volume.</p>	<p>No user action is required.</p> <p>Although you can safely access a mirrored volume while resynchronization is in progress, you should avoid making configuration changes (such as breaking a mirror) during resynchronization.</p>
Regenerating	<p>Data and parity are regenerated for a RAID-5 volume. Both dynamic and basic RAID-5 volumes display the Regenerating status.</p> <p>When regeneration is complete, the RAID-5 volume's status returns to Healthy.</p>	<p>No user action is required.</p> <p>You can safely access a RAID-5 volume while data and parity regeneration is in progress.</p>
Degraded	<p>The Degraded status applies only to mirrored or RAID-5 volumes on basic or dynamic disks. There are three situations where Degraded status can occur:</p>	<p>Take the following actions to resolve these situations:</p>
	<p>Situation 1</p> <p>The data on the volume may no longer be fault tolerant because one of the underlying disks is not online. A RAID-5 volume loses redundancy if one disk is offline or failed. A mirrored volume can lose redundancy if one of the disks containing one of its subdisks fails.</p>	<p>Situation 1</p> <p>You can continue to access the volume using the remaining online disks, but you should repair the volume as soon as possible.</p> <p>It is important to repair the volume because if another disk that contains the volume fails, you lose the volume and its data.</p> <p>See "Bring a dynamic volume back to a healthy state" on page 533.</p> <p>See "Bring a basic volume back to a healthy state" on page 532.</p>

Table 20-2 Volume Status Descriptions (*continued*)

Status	Meaning	Action Required
	<p>Situation 2</p> <p>A Degraded status also displays if a disk involving a RAID-5 or mirrored volume was physically moved.</p>	<p>Situation 2</p> <p>To correct the problem, you must move all the disks that contain the volume to the new location or return the moved disk or disks to the original location.</p>
	<p>Situation 3</p> <p>The data on the volume is no longer fault tolerant, and I/O errors have been detected on the underlying disk. If an I/O error is detected on any part of a disk, all volumes on the disk display the (At Risk) status. Only dynamic mirrored or RAID-5 volumes display the Degraded (At Risk) status.</p>	<p>Situation 3</p> <p>To return the underlying disk to the Online status, reactivate the disk (using the Reactivate Disk command). Once the disk is returned to the Online status, the volume status should change to Degraded.</p> <p>Take further action as necessary to bring the volume back to a healthy state.</p> <p>See “Bring a dynamic volume back to a healthy state” on page 533.</p>
Failed	<p>A failed volume starts automatically. An error icon appears on the failed volume. Both dynamic and basic volumes display the Failed status. There are two situations where Failed status can occur:</p>	<p>Take the following actions to resolve these situations:</p>
	<p>Situation 1</p> <p>A volume that spans two or more disks has one or more disks fail. A striped volume, simple volume, spanned volume, or extended partition fails with one disk failure. A RAID-5 volume fails with two disk failures. A mirrored or mirrored-striped volume fails when the disks containing all the mirrors in the volume fail.</p>	<p>Situation 1</p> <p>Replace or repair the failed or malfunctioning disk or disks.</p>

Table 20-2 Volume Status Descriptions (*continued*)




Status	Meaning	Action Required
	<p>Situation 2</p> <p>A disk or disks involving a volume that spans multiple disks is moved to another computer.</p>	<p>Situation 2</p> <p>Move all the disks that contain subdisks of the volume, or return the moved disk or disks to the original location.</p>
Formatting	The volume is formatted using the specifications you chose for formatting.	No user action is required.
Stopped	The volume is in a dynamic disk group that is not imported.	Import the dynamic disk group containing the volume.
Missing	The volume status is Missing if any of the subdisks of the volume are on the disks that show an Offline status.	Reactivate the offline disks and Rescan. If the volume status changes to Stopped or Failed, reactivate the volume.
Failing	Failing is a secondary message shown in parentheses after a volume status. Failing means Storage Foundation encountered some I/O errors on at least one disk that contains subdisks of the volume; however, these errors did not compromise the data on the specific volume. Failing sends the message that the disk integrity is deteriorating. When the volume status is Degraded (At Risk), the underlying disk's status is usually Online (Failing).	<p>Determine which disk is failing, and take corrective action.</p> <p>See "Disk status descriptions" on page 521.</p>

SFW error symbols

When a problem occurs in a computer's storage subsystem, Storage Foundation alerts the user with error messages and error symbols that are placed on top of the disk icons or volume icons to show the source of the problem.

The following table lists these error symbols, explains their meaning, and gives examples of their usage.

Table 20-3 Storage Foundation Error Symbols

Symbol	Meaning	Example
	Warning. The yellow caution symbol indicates there is a potential problem but the system can still function normally.	Capacity warning, Degraded warning.
	Informational. The blue information symbol indicates that an object is not completely accessible or fully functional; it also indicates various informational events. This state is often caused by a user operation. No degradation or loss of data; the system can still function normally.	Disk group is deported, disks do not have signatures, disk group has some legacy objects (such as DCO, snapshotted volumes) that would not work until the disk group is upgraded, shared volume is not online on the node, server started, device arrival, etc.
	Failure. The X symbol on a red circle indicates that a serious storage failure has occurred. Users need to look at the source of the problem and fix it as soon as possible. Any attempt to access the object results in an error.	Failed volume, Missing disk.

Resolving common problem situations

This section provides information on how to resolve the most common problem situations that users encounter when managing their storage with Storage Foundation.

See [“Bring an offline dynamic disk back to an imported state”](#) on page 529.

See [“Bring a basic disk back to an online state”](#) on page 530.

See [“Remove a disk from the computer”](#) on page 531.

See [“Bring a foreign disk back to an online state”](#) on page 531.

See [“Bring a basic volume back to a healthy state”](#) on page 532.

See [“Bring a dynamic volume back to a healthy state”](#) on page 533.

See [“Repair a volume with degraded data after moving disks between computers”](#) on page 534.

See [“Deal with a provider error on startup”](#) on page 535.

Bring an offline dynamic disk back to an imported state

An Offline dynamic disk may be corrupted or intermittently unavailable.

To bring an offline dynamic disk back to an imported state

- 1 Repair any disk or controller problems, and make sure that the disk is turned on, plugged in, and attached to the computer.
- 2 Use the **Rescan** command to rescan all the devices on the SCSI bus to bring the disk back online.

Select **Rescan** from the **Actions** menu, or right-click the StorageAgent node in the tree view to get a context menu and select **Rescan** from that menu. If you have a lot of devices on the computer’s SCSI bus, the rescan of the bus may take some time.

See [“Rescan command”](#) on page 536.

If one of the disks has failed and you have a mirrored or RAID-5 volume, the repair involves recreating a part of the volume in a different location.

For more information:

See [“Repair volume command for dynamic RAID-5 volumes”](#) on page 542.

See [“Repair volume command for dynamic mirrored volumes”](#) on page 542.

- 3 If the disk does not come back after doing a rescan, select the disk and use the **Reactivate Disk** command to manually bring the disk back online.

Right-click the disk’s label in the tree view or the Disk View tab view to bring up the disk’s context menu and select the **Reactivate Disk** command from the menu.

See [“Reactivate disk command”](#) on page 540.

Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Disk** command.

If reactivating the disk does not change its status, something must be wrong with the disk or its connections.

- 4 If a disk comes back online after reactivating, check to see whether its volumes are healthy. If not, try the **Reactivate Volume** command on its volumes.

See [“Reactivate volume command”](#) on page 541.

Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Volume** command.

- 5 Run **Chkdsk.exe** to make sure that the underlying data on the disk is not corrupted.

To run **Chkdsk**, open a command prompt window and enter the following command:

```
chkdsk x: /f
```

where *x* is the drive letter of the volume you want to check. The **/f** option tells **Chkdsk** to fix any errors it finds. If the **/f** option is omitted, **Chkdsk** operates in a read-only mode.

Although **Chkdsk** cleans up the file system structure, there may be invalid data on the disk if users are working when a problem occurred on the disk. It is highly recommended that you run utilities to check the integrity of the data. If the data is corrupted, you need to replace it from backup storage.

Bring a basic disk back to an online state

If a basic disk is corrupted or unavailable, it does not show in the VEA GUI.

To bring a basic disk back to an online state

- 1 Repair any disk or controller problems, and make sure that the disk is turned on, plugged in, and attached to the computer.
- 2 Use the **Rescan** command to rescan all the devices on the SCSI bus to bring the disk back online.

Select **Rescan** from the **Actions** menu, or right-click the StorageAgent node in the tree view to get a context menu and select **Rescan** from that menu. If you have a lot of devices on the computer’s SCSI bus, the rescan of the bus may take some time.

See [“Rescan command”](#) on page 536.

- 3 If a disk comes back online after rescanning, check to see whether its volumes are healthy.

If they are not, you need to bring the volumes back to a healthy state.

See [“Bring a basic volume back to a healthy state”](#) on page 532.

- 4 Run **Chkdsk.exe** to make sure that the underlying data on the disk is not corrupted.

Even if the disk and volumes come back online, it is important to check whether the underlying data is intact.

To run **Chkdsk**, open a command prompt window and enter the following command:

```
chkdsk x: /f
```

where *x* is the drive letter of the volume you want to check. The **/f** option tells **Chkdsk** to fix any errors it finds. If the **/f** option is omitted, **Chkdsk** operate in a read-only mode.

If the data is corrupted, you need to replace it with data from backup storage.

Remove a disk from the computer

As part of troubleshooting you may need to remove a basic or dynamic disk from the computer. This information is covered in the topics on disk tasks.

See [“Remove a disk from the computer”](#) on page 103.

Bring a foreign disk back to an online state

To bring a foreign disk back to an online state depends on the original context of the disk.

Note: Because a volume can span more than one disk (such as with a mirrored or RAID-5 volume), it is important that you first verify your disk configurations and then move all the disks that are part of the volume. If you do not move all the disks, the volume has a Degraded or a Failed error condition.

The following describe the various contexts:

- Use Import Dynamic Disk Group to bring a Foreign disk online if the disk was part of a dynamic disk group that was created on another computer and moved to the current computer.

Note: Microsoft Disk Management Disk Groups do not support the Import Dynamic Disk Group command.

To add a foreign disk that was created on another computer

- 1** Right-click on the disk and select **Import Dynamic Disk Group**.

A dialog box comes up that displays the name of the dynamic disk group.

- 2** Indicate the name for the dynamic disk group as follows:

- If you want to leave the name as is, click **OK**.
- If you want a new name for the dynamic disk group, type a new name for the group in the **New name** entry box and then click **OK**.

- 3** If you import a dynamic disk group from a different system, then click the check box to clear the host ID of the other system.

The disk group then imports. All existing volumes on the disk are visible and accessible.

More information on importing and deporting dynamic disk groups is available.

See [“Importing and deporting dynamic disk groups”](#) on page 157.

- Use Import Dynamic Disk Group if the Foreign disk has a secondary dynamic disk group (that is, a dynamic disk group other than the dynamic disk group containing the computer’s boot or system disk) and you have switched between operating systems on a dual-boot computer.

When you have one or more secondary dynamic disk groups on a disk in a dual-boot environment, the disk is marked as Foreign when you switch between operating systems, because the secondary disk groups are not automatically imported. A shared primary dynamic disk group on a disk in that situation is automatically imported.

- Use Merge Foreign Disk if the disk was originally created on the current computer but was removed and now you have reattached it to the current computer and want to restore it to its status as a member its original dynamic disk group.

See [“Merge foreign disk command”](#) on page 538.

Bring a basic volume back to a healthy state

The following steps bring back a basic volume to a healthy state.

To bring a basic volume back to a healthy state

- 1 Repair any disk or controller problems, and make sure that the disk is turned on, plugged in, and attached to the computer.
- 2 Use the **Rescan** command to rescan all the devices on the SCSI bus to bring the disk that the volume belongs to back online.

Select **Rescan** from the **Actions** menu, or right-click the StorageAgent node in the tree view and select **Rescan** from the context menu. If you have a lot of devices on the computer's SCSI bus, the rescan of the bus may take some time.

See ["Rescan command"](#) on page 536.

Bring a dynamic volume back to a healthy state

The following steps bring back a dynamic volume to a healthy state.

To bring a dynamic volume back to a healthy state

- 1 First you need to attempt to bring the disk or disks to which the volume belongs to Online status by doing a **Rescan** and, if necessary, a **Reactivate Disk**.

Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Disk** command.

See “[Rescan command](#)” on page 536.

See “[Reactivate disk command](#)” on page 540.

If one of the disks has failed and you have a mirrored or RAID-5 volume, the repair involves recreating part of the volume in a different location.

See “[Repair volume command for dynamic RAID-5 volumes](#)” on page 542.

See “[Repair volume command for dynamic mirrored volumes](#)” on page 542.

- 2 If a disk comes back online after reactivating, check to see whether its volumes are healthy. If not, try the **Reactivate Volume** command on its volumes.

Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Volume** command.

See “[Reactivate volume command](#)” on page 541.

- 3 Run **Chkdsk.exe** to make sure that the underlying file system structure is intact.

To run **Chkdsk**, open a command prompt window and enter the following command:

```
chkdsk x: /f
```

where *x* is the drive letter of the volume you want to check. The **/f** option tells **Chkdsk** to fix any errors it finds. If the **/f** option is omitted, **Chkdsk** operates in a read-only mode.

Although **Chkdsk** cleans up the file system structure, the disk may have invalid data if users are working when a problem occurred on the disk. It is highly recommended that you run utilities to check the integrity of the data. If the data is corrupted, you need to replace it from backup storage.

Repair a volume with degraded data after moving disks between computers

This section gives steps you can take if you used the **Deport Dynamic Disk Group** and **Import Dynamic Disk Group** commands to move disks between computers and the disks contain mirrored or RAID-5 dynamic volumes with degraded data.

To resolve the problem:

- 1 Deport the disks on the computer they were moved to and then physically move the disks back to the computer where they originated.
- 2 Use **Rescan** to make sure all the disks are installed correctly.
 If the volume showed Degraded status before the disks were moved, the volume continues to show a Degraded status when the disks are moved back.
- 3 Make sure the status of the disk that contained the degraded mirror or parity information is not Offline.
 If the status is Offline, check for any hardware problems, and reconnect the disk, if necessary.
- 4 Use the **Reactivate Disk** command to bring the disk back online.
 Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Disk** command.
 If the hardware problem is corrected, the disk shows a Healthy status, any mirrored volumes on the disk are resynchronized, and any RAID-5 volumes generate parity.
- 5 If any volume still shows Degraded status, use the **Reactivate Volume** command on that volume.
 Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Volume** command.
 Storage Foundation attempts to bring the volume back online. You may see a warning message indicating that Storage Foundation may not be able to bring back all of the data on the volume.
 If Storage Foundation successfully brings the volume back online, the status of the volume becomes Healthy.
- 6 Now you can deport the dynamic disk group and move all the dynamic disk group's disks to the second computer.
 Be sure to move all the disks that are involved in the disk group at the same time to ensure that your volumes have the Healthy status on the second computer.

Deal with a provider error on startup

In Storage Foundation, providers are similar to drivers. Each provider manages a specific hardware or software storage component. For example, there is a disk provider that manages all disks that the Windows operating system sees as disks.

The providers discover the existing physical and logical entities and store that information in the Storage Foundation distributed database.

If you receive an error at Storage Foundation startup that a provider failed to load, you can check the provider status by right-clicking the managed server node in the Veritas Enterprise Administrator tree view and selecting Properties from the context menu that appears.

The top section of the window displays the loaded providers. The bottom section of the window displays any providers that failed to load. Normally, the bottom section is blank.

If a certain provider failed to load at SFW startup, the features that the provider supplies does not work in SFW. It is necessary to determine why the provider did not load and start the application again. For assistance, contact Veritas Technical Support.

Commands or procedures used in troubleshooting and recovery

This section covers ways to perform troubleshooting and recovery.

Refresh command

If you do not believe that anything is wrong with the disks or volumes but the VEA GUI has not updated a recent change, try the Refresh command. Refresh refreshes drive letter, file system, volume, and removable-media information on the current computer. It also checks to see whether previously unreadable volumes are now readable. It may not pick up the disk changes that occurred since the last restart or rescan if there was no I/O occurring on the changed disk.

Select **Refresh** from the Actions or View menu of the VEA, or right-click the **StorageAgent** node in the tree view to get a context menu and select **Refresh** from that menu.

Note: The Refresh command is only effective on the StorageAgent node and any of its sub-nodes.

Rescan command

The Rescan command rescans the SCSI bus for disk changes. It also does the equivalent of the Refresh command, updating information on drive letter, file system, volume, and removable media.

It is recommended that you use Rescan every time you make disk changes, such as removing or adding a disk. Rescanning can take several minutes, depending on the number of devices on the SCSI bus.

From the toolbar, select Actions, then Rescan.

A progress bar showing the percentage of completion for the rescan process is available by clicking on the **Tasks** tab located in the lower pane. When Rescan is finished, you should see accurate information about your system.

If error flags remain, you may need to reactivate the disks or volumes.

Note: The Rescan command is only effective on the StorageAgent node and any of its sub-nodes.

Additional important information about the rescan command

This section gives additional information on how a rescan affects redundant and nonredundant volumes.

The details are as follows:

- Redundant volumes (Mirrored and RAID-5)

If a RAID-5 or a mirrored volume is in a Degraded state, doing a rescan also automatically triggers an attempt to reactivate the disks and resynchronize the volume on the disks. This is done in the background and is the reason why it is not always necessary to invoke the Reactivate Disk and Reactivate Volume commands.
- Nonredundant volumes (Simple, spanned, and striped)

If you have a simple, spanned, or striped volume that has failed, a rescan does not automatically attempt to reactivate the disk and the volumes. The data on these failed volumes may already be corrupted.

For a failed simple, spanned, or striped volume, you must manually reactivate the disk(s) and the volume. The error messages you receive may give you a hint of the reason for the problem. After manually reactivating the disk(s) and volume, the disks may come back as Online and the volume as Healthy, but you cannot assume that the underlying data has not been affected.

In this situation, you should also run Chkdsk. Although Chkdsk cleans up the file system structure, the disk may have invalid data if users are working when a problem occurred on the disk. It is highly recommended that you run utilities to check the integrity of the data. If the data is corrupted, you need to replace it from backup storage.

Note: If you have mirrored or RAID-5 volumes with a Failed state, a rescan does not automatically reactivate the disks and resynchronize the volumes on the disks. Once a redundant volume has failed, it is subject to the same conditions as a failed nonredundant volume.

Replace disk command

This command lets you replace a failed disk with an empty basic disk. The volume configuration is recreated on the new disk. The contents of the nonredundant volumes are not guaranteed. Redundant volumes are automatically resynchronized. This command can be performed only for a disk that is designated as "Missing."

See ["Replace disk"](#) on page 106.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Replace Disk command.

Merge foreign disk command

If you remove a disk from the server and also remove it in Storage Foundation and then decide to reattach it to the server as a member of the same dynamic disk group, you need to use the Merge Foreign Disk command. This command reinstates the disk to its old status as a member of its former dynamic disk group on that server.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Merge Foreign Disk command.

You also need to use this command if you remove a disk without taking its disk group offline, move the disk to another server, and then attempt to reinstall it in the original server. The command is necessary in this case because the disk has the other server's disk group ID.

In Storage Foundation, this command does not need to be used very often, because if a disk is removed from a computer that has subdisks from a redundant volume (such as a RAID-5 volume or a mirrored volume), the Hot Relocation command automatically rebuilds the missing subdisk on the volume. However, if the hot relocation feature is turned off, you may need this command.

See ["Hot relocation"](#) on page 216.

If you have a dynamic disk group with one disk marked as Missing and you revert the remaining disks in the group to basic, the dynamic disk group no longer exists,

and the disk marked as Missing also disappears from the VEA GUI. This happens because the information about a dynamic disk group is stored in a database on each of the disks that reside in the group. The dynamic disk group information is no longer available from the disks that were reverted to basic, and the Missing disk is not connected. If you reconnect the Missing disk and rescan, it shows up as Foreign. Now its database is available, but the remaining disks in the group are gone.

More information about removing a dynamic disk from the computer is available.

See [“Remove a disk from the computer”](#) on page 531.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support hot relocation.

Note: Dynamic disks belonging to a SAN Disk Group do not support hot relocation.

To use the merge foreign disk command

- 1** Reattach the disk to the original server.
- 2** In the VEA console, do a rescan (select **Rescan** from the **Actions** menu). The disk is displayed in the tree with a red X. Its dynamic disk group displays as "Unknown Group."
- 3** Right-click the disk label in the tree view to bring up the disk context menu.
- 4** Select **Merge Foreign Disk** from the disk context menu.
- 5** After reading the first screen of the Merge Foreign Disk wizard, click **Next** to continue.
- 6** In the Merge Foreign Disk wizard screen that appears, select the disk that you want to merge.

Using the **Add** option, move it from the left pane of the screen to the right pane and then click **Next** to continue.

- 7 In the next screen, click **Next** to continue if the data condition on the disk is **Healthy**.

The final screen of the Merge Foreign Disk wizard appears.

The volume condition types are shown below:

Failed Volume is not functional. Volume will be merged, may be unusable.

Failed Redundancy Volume contents are not current. Volume will be merged.

Missing Volume does not have a subdisk on merged disks.

If the volume condition type is **Failed**, there is no guarantee that the data is intact (though it may be).

See [“Bring a dynamic volume back to a healthy state”](#) on page 533.

If the disk is **Missing**, make sure that it is properly connected.

- 8 Click **Finish** to complete the process of merging the Foreign disk into the server.

The merged disk should now appear as it was before it was removed from the server. It should now be shown as being in its former dynamic disk group.

- 9 If the disk still has an error symbol, right-click the disk and click **Reactivate Disk** from the context menu that appears.

The disk now should appear as normal.

Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the **Reactivate Disk** command.

Reactivate disk command

Occasionally, Rescan does not clear error flags on a dynamic disk. When this happens, you can use Reactivate Disk to clear these flags. This command does a manual restart on a disk. Dynamic disks that are marked as **Missing** or **Offline** can be reactivated. After reactivating, disks should be marked as **Online** if they are attached and are not in a **Failed** condition.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Reactivate Disk command.

To clear error flags using the reactivate disk command

- 1 Right-click the disk with the error flags and select **Reactivate Disk** from the context menu that appears.

A dialog box appears with a confirmation message.

- 2 Click **Yes** to reactivate the disk.

The disk should be marked Online after it is reactivated unless there is a mechanical or other serious problem with the disk.

See [“Additional important information about the rescan command”](#) on page 537.

Reactivate volume command

If you have errors on a dynamic volume, you must try to bring any disks on which it resides back online by first using Rescan and then, if that does not work, Reactivate Disk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Reactivate Disk command.

Once the disk or disks come back online, if the volume does not return to a Healthy state, use the Reactivate Volume command.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the Reactivate Volume command.

Right-click on the volume to bring up the context menu, and then click **Reactivate Volume**.

This can be necessary if any of the mirrors or plexes of a mirrored volume are still not healthy. The same situation applies to a RAID-5 volume; you may need to use Reactivate Volume to allow the volume to regenerate.

If the underlying disks for a volume are sound, the volume most likely comes back to a Healthy state, but the data on it may be corrupted or stale. It is recommended that you run Chkdsk.exe before using the volume. If Chkdsk fails or finds an excessive number of errors, you may need to reformat the volume and restore its contents from backup.

The difference between the Reactivate Volume and Repair Volume commands is that Reactivate Volume resynchronizes the volume to bring it to a Healthy state in its existing location, while Repair Volume removes the damaged section of the volume and recreates it in another location on a healthy dynamic disk.

See [“Additional important information about the rescan command”](#) on page 537.

Repair volume command for dynamic RAID-5 volumes

You can repair a RAID-5 volume if its status is Degraded and if there is enough unallocated space available on other dynamic disks to recreate the degraded subdisks of the volume. To avoid data loss, you should attempt to repair the volume as soon as possible.

Note: After a disk fails, you may need to rescan before this menu option is available.

The difference between the Reactivate Volume and Repair Volume commands is that Reactivate Volume resynchronizes the volume to bring it to a Healthy state in its existing location, while Repair Volume removes the damaged section of the volume and recreates it in another location on a healthy dynamic disk.

To repair a dynamic RAID-5 volume using the repair volume command

- 1 Right-click the degraded volume. Select **Repair Volume** from the context menu that appears.

The Repair Volume dialog box appears.
- 2 The default setting is for Storage Foundation to automatically assign the destination disks. To manually select the disks, click the **Manually assign destination disks** radio button and then select the disks you want assigned as destination disks. You may also check **Disable Track Alignment** to disable track alignment for the recreated volume.
- 3 When you are satisfied with your selection, click **OK**.

Repair volume command for dynamic mirrored volumes

When a disk on which a mirrored volume resides fails, the volume displays a Degraded status. The disk’s name is changed to Missing Disk, and an icon (X) appears on the Missing Disk icon. The status of the disk is Offline.

The difference between the Reactivate Volume and Repair Volume commands is that Reactivate Volume resynchronizes the volume to bring it to a Healthy state in its existing location, while Repair Volume removes the damaged section of the volume and recreates it in another location on a healthy dynamic disk.

Using the repair volume command

The following describes how to repair a mirrored volume using the repair volume command.

To repair a mirrored volume using the repair volume command

- 1 Right-click the degraded volume. Select **Repair Volume** from the context menu that appears.

The Repair Volume dialog box appears.

- 2 Select the mirror or mirrors to repair by clicking in the corresponding checkboxes. You may also check **Disable Track Alignment** to disable track alignment for the recreated mirrors.

- 3 Click **OK**.

New mirrors are created on available disk space on other dynamic disks.

Repairing a volume and manually selecting target disks

The following describes how to repair a volume and manually select target disks for new mirrors.

To repair a volume and manually select the target disks for the new mirrors

- 1 Right-click the degraded volume. Select **Mirror** from the context menu, then **Remove** from the submenu.

The **Remove Mirror** dialog box appears.

- 2 Remove the degraded mirror.

See [“Remove mirror”](#) on page 211.

Note that the mirrors on preserved disks (in the right pane) are not removed.

- 3 Right-click the volume again. Select **Mirror** from the context menu, then **Add** from the submenu.

- 4 Manually select the target disk to add the mirror to and then click **OK**.

See [“Add a mirror to a volume”](#) on page 208.

Starting and stopping the Storage Foundation Service

It is useful to know how to start and stop the Storage Foundation service when you are troubleshooting. For example, if Storage Foundation stops running on the server, you can try restarting the service instead of restarting. Sometimes stopping the service and restarting it again can resolve temporary problems. The Storage Foundation service is also referred to as the `vxsvc`.

Note: For the current release, to start or stop the Storage Foundation service you start or stop the `vxsvc` service (or the Veritas Storage Agent in the Windows Services GUI).

The following describes how to start and stop the Storage Foundation service:

- Start the Storage Foundation service.
 Bring up the command window and enter the following command:

```
net start vxsvc
```

Note: If you restart, the service is automatically restarted.

- Stop the Storage Foundation service.
 Bring up the command window and enter the following command:

```
net stop vxsvc
```

Accessing the CLI history

The history of CLI commands issued on a system and the system responses is contained in the command line log, `vxcli.log`. It is typically located at `C:\Program Files\Veritas\Veritas Volume Manager\logs`. This history is useful when trying to understand the problems that may arise from running or creating CLI scripts.

Each entry in the log is prefixed with the date and time of the command or system response, followed by the process identification and thread identification. Each CLI command is a separate entry in the log and each is identified by "----- Command : ". Each system response is also a separate entry in the log and follows the CLI command.

Note: Help messages are not recorded in the command line log.

Additional troubleshooting issues

This section describes additional troubleshooting issues.

See [“Disk issues”](#) on page 545.

See [“Volume issues”](#) on page 546.

See [“Disk group issues”](#) on page 549.

See [“Connection issues”](#) on page 552.

See [“Issues related to boot or restart”](#) on page 553.

See [“Cluster issues”](#) on page 555.

See [“Dynamic Multi-Pathing issues”](#) on page 558.

See [“vxsnap issues”](#) on page 558.

See [“Other issues”](#) on page 560.

Disk issues

This section contains information about troubleshooting issues with disks.

See [“Disk type shows no signature”](#) on page 545.

See [“A disk is marked as foreign”](#) on page 545.

See [“Error when upgrading a basic disk to dynamic”](#) on page 545.

Disk type shows no signature

If the type of disk shows No Signature, you need to write a signature to the disk. When installing a new disk, the software must write a signature to the disk that prepares it for use. This signature is not written automatically, in case that disk has been imported from another operating system, and the configuration information needs to be kept intact.

Check the following:

- To write the configuration data to the disk, right-click on the disk under the Disks node and choose Write Signature.

A disk is marked as foreign

This section contains information about foreign disks.

See [“Bring a foreign disk back to an online state”](#) on page 531.

See [“Disk status descriptions”](#) on page 521.

Error when upgrading a basic disk to dynamic

This topic describes known problems when upgrading a basic disk to dynamic.

Basic disk has more than 32 logical drives

If you attempt to upgrade a basic disk that has more than 32 logical drives in its extended partition, it does not upgrade to a dynamic disk. You get the error message “Disk cannot be partitioned after it has been upgraded.” There is a limit of 16 logical drives for a basic disk to be upgraded to a dynamic disk.

Basic disk has an "Active" volume with no system files

If you mark a basic mirrored volume active that does not contain the current system files, the computer is not able to restart. A basic disk that contains an unbootable active volume cannot be upgraded to dynamic.

You receive an error about disk geometry

If you receive an error about disk geometry when you try to upgrade a basic disk to a dynamic disk and the disk does not upgrade, the geometry on the disk is inconsistent with the disk geometry used on the current Windows operating system. If the disk is not upgraded, Storage Foundation maintains the disk with its current partition structure; but if the partition structure is changed in any way, the disk partitions, and their data become inaccessible.

The following describes the fix for the situation:

- Back up all your data from the affected disk.
- Delete all partitions on the disk.
- Upgrade the disk to dynamic.
- Create volumes on the disk.
- Restore your data to the disk.

It is very likely that you have a disk that was created under NT 4.0 with partitions on it and this disk has a different geometry because it was attached to a secondary host bus adapter with the BIOS turned off and mapping features turned off.

Once you have done the above steps, the disk has the proper geometry, and can still be connected to the secondary host adapter with the BIOS and mapping features turned off. The geometry problem was the result of the change from NT 4.0 to Windows 2000. It is a known problem and is discussed in more detail on the Microsoft Web site.

Volume issues

This section contains information about troubleshooting issues with volumes.

See [“After a failover, VEA sometimes does not show the drive letter or mounted folder paths of a successfully-mounted volume”](#) on page 547.

See [“Cannot create a RAID-5 volume”](#) on page 548.

See [“Cannot create a mirror”](#) on page 548.

See [“Cannot extend a volume”](#) on page 548.

See [“Cannot extend a dynamic boot or system volume”](#) on page 548.

See [“When creating a spanned volume over multiple disks within a disk group, you cannot customize the size of subdisks on each disk”](#) on page 549.

See [“An attempt to mirror a boot or system volume fails or has an error message”](#) on page 549.

After a failover, VEA sometimes does not show the drive letter or mounted folder paths of a successfully-mounted volume

This is related to failover, when VEA sometimes does not show the drive letter or mounted folder paths of a volume even though the volume is successfully mounted with the expected drive letter or folder paths. During a failover, when a disk group gets imported, SFW mounts all volumes of the disk group by querying the mount points using Microsoft API `GetVolumePathNamesForVolumeName()`. Sometimes, this API fails to return the correct drive letter or mounted folder paths because of which VEA fails to update the same.

To resolve this, the operation needs to be retried after a few milliseconds so that the Microsoft API `GetVolumePathNamesForVolumeName()` returns correct information, but it also affects the performance.

As a workaround, you need to enable the retry logic by creating the registry keys that are mentioned below.

Note: Please note that using this workaround has a performance impact on the disk group deport operation performed while the service group is offline. It affects the performance by a maximum of $n/2$ seconds, where "n" is the number of volumes in the disk group.

A retry logic has been added for the `GetVolumePathNamesForVolumeName()` API so that it retries the operation in case the returned mount path is empty. It retries after every 100 milliseconds for "n" number of attempts, where "n" is configurable through registry.

To enable this retry logic, create the registry value named "RetryEnumMountPoint" with the type `REG_DWORD` under the following registry key and set its value to "1":

```
HKKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager
```

To specify the number of retry attempts, create a registry value named "RetryEnumMPAttempts" with the type `REG_DWORD` under the following registry key and set its value to "5". You can change its value as needed.

```
HKKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager
```

Cannot create a RAID-5 volume

This section contains information about not being able to create a RAID-5 volume.

Check the following:

- Adequate unallocated space on three or more disks? You must have at least three disks to create a RAID-5 volume or four disks to create a RAID-5 volume with a log.

Cannot create a mirror

This section contains information about not being able to create a mirror.

Check the following:

- Adequate unallocated space on two or more dynamic disks? You must have two or more disks to create a mirrored volume.

See [“An attempt to mirror a boot or system volume fails or has an error message”](#) on page 549.

Cannot extend a volume

One possible reason for not being able to extend a volume is that the RAID-5 or DRL log that was associated with the volume cannot be found, possibly because a disk has failed, or is offline. The workaround is to delete the log from the failed disk or bring the missing disk back online, and then you should be able to extend the volume. If your volume has a boot or system disk, see the next section.

Cannot extend a dynamic boot or system volume

The following are possible reasons for this issue:

- A boot or system volume cannot be extended unless there is enough contiguous space after the volume for the desired extension. If you have another volume on the disk, you cannot use space before that volume and after that volume for the extension.
- If the extension of a dynamic boot or system volume requires adjustment of an extended partition boundary, then the extension is not allowed. The error message "Requested operation not supported" is returned.
- A boot or system disk volume that was originally created on a logical drive and later upgraded to a Storage Foundation dynamic volume cannot be extended in Storage Foundation.

When creating a spanned volume over multiple disks within a disk group, you cannot customize the size of subdisks on each disk

When you create a spanned volume over multiple disks, Storage Foundation does not let you use the Custom mode to create subdisks of a specified size on the disks in the volume. Storage Foundation automatically uses all the available disk space in the first disk before moving to the second disk, and likewise does the same with the second disk before moving to the third, and so on.

To work around this problem, create a simple volume on the first disk with a specified size (the size is a fraction of the total size needed). Once this is created, use the Expand Volume command with the Custom mode to specify additional space on another disk. Repeat for the total number of disks being used.

This technique lets you create a spanned volume with specific subdisk sizes on each disk in the group.

An attempt to mirror a boot or system volume fails or has an error message

The following are possible causes for this problem:

- **The Only Available Dynamic Disk Large Enough for the Mirror Has More than Two Retained Partitions**
 If you try to add a mirror to a system or boot disk but the only dynamic disks that have enough space contain more than two retained partitions, the Add Mirror operation fails. The error message says "insufficient disk space," even though there is plenty of space on the disk.
- **No Available Dynamic Disk Has the Same Offset as the Original Boot or System Disk**
 To boot from a mirrored system volume, its offset from the beginning of the disk must be in the same location as that of the original system volume's offset. If you try to add a mirror to a system volume where there are no dynamic disks that can have the offset in the same location, you receive a message to that effect and be asked whether you want to continue. You can choose to continue, so that a mirror at a different offset is created. The mirror correctly mirrors the data, but you cannot boot from the mirror.

Disk group issues

This section contains information about troubleshooting issues with disk groups.

See ["Unknown group appears after upgrading a basic disk to dynamic and immediately deporting its dynamic disk group"](#) on page 551.

See [“Cannot use SFW disk groups in disk management after uninstalling InfoScale Storage management software”](#) on page 551.

See [“After uninstalling and reinstalling InfoScale Storage management software, the private dynamic disk group protection is removed”](#) on page 551.

See [“Cannot import a cluster dynamic disk group or a secondary disk group with private dynamic disk group protection when SCSI reservations have not been released”](#) on page 552.

See [“Dynamic Disk Group Split and Join troubleshooting tips”](#) on page 285.

Unable to fail over an existing Volume Manager Shared Volume resource to a newly-added Microsoft failover cluster node

Note: This issue may also occur for the Volume Manager Disk Group (VMDg) and Replicated Volume Group (RVG) resources. The following information is applicable in case of those two resources as well.

On Windows Server 2008 R2 operating systems, this issue occurs while failing over an existing Volume Manager Shared Volume resource of a cluster-shared disk group (CSDG) to a node that is newly added to the Microsoft failover cluster and then to the CVM cluster. The operation fails with the following error message:

"Error Code: 0x80071398" The operation failed because either the specified cluster node is not the owner of the group, or the node is not a possible owner of the group.

This is a Microsoft failover cluster issue.

Workaround: To resolve this issue, you need to manually add the new node to the Volume Manager Shared Volume resource using Microsoft’s Failover Cluster Manager snap-in, and then retry the failover. In the snap-in GUI, right-click the Volume Manager Shared Volume resource and select **Properties**. In the resource properties window, go to the **Advanced Policies** tab, select the newly-added node, and then click **OK** to add the node as a possible owner of the resource.

Sometimes, creating dynamic disk group operation fails even if disk is connected to a shared bus

This happens while creating a cluster dynamic disk group. In some cases, this operation fails if the disk (to which the disk group belongs) contains an OEM partition even if it's connected to a shared bus. You can resolve this issue by deleting the OEM partition and restarting the Veritas Enterprise Administrator Service (VxSVC) service.

Unknown group appears after upgrading a basic disk to dynamic and immediately deporting its dynamic disk group

Upgrading a basic disk and immediately deporting its dynamic disk group occasionally results in the appearance of a dynamic disk group named "Unknown." Refreshing the display or trying to import the deported dynamic disk group removes the original group from the display and, in the case of import, generates an error that the disk cannot be found.

Do not attempt any other operations on these disks from within Storage Foundation. Doing so can result in a loss of data. To recover the dynamic disk group and its contents requires a restart of the computer. After restarting, the dynamic disk group displays correctly as a deported group as "Offline, Foreign." Then you can import it without any problems.

Cannot use SFW disk groups in disk management after uninstalling InfoScale Storage management software

After uninstalling the InfoScale Storage management software, the existing disk group may not be able to be imported and used in Disk Management.

When you uninstall InfoScale Storage management software, Disk Management automatically imports only primary disk groups.

If, in Storage Foundation, you do not have a primary disk group because your system or boot disk is not encapsulated, and then you uninstall InfoScale Storage management software, Disk Management is unable to import the disk group or groups, because it cannot import a secondary disk group as a primary disk group. Without a primary disk group, you cannot merge the foreign disk group.

The workaround is to create a new dynamic disk group in Disk Management and then the Foreign disk group(s) can be merged into the dynamic disk group.

After uninstalling and reinstalling InfoScale Storage management software, the private dynamic disk group protection is removed

If you uninstall InfoScale Storage management software and then reinstall it again (either through an upgrade or for other purposes), the private dynamic disk group protection is removed during the uninstall process. If you then reinstall InfoScale Storage management software and want to continue the private dynamic disk group protection for the disk groups that had it previously, add the feature back to the disk groups through the **Add Dynamic Disk Group Protection** command. Private dynamic disk group information is stored in the registry; and, by necessity, the registry entries are removed during an uninstall.

Cannot import a cluster dynamic disk group or a secondary disk group with private dynamic disk group protection when SCSI reservations have not been released

Importing a cluster dynamic disk group or a secondary disk group with private dynamic disk group protection may not be possible when a SFW server on a shared bus had failed to release SCSI reservations when the cluster dynamic disk group or secondary disk group was deported. Performing the clear SCSI reservation operation releases these reservations and allows importing the cluster dynamic disk group or secondary disk group.

To clear SCSI reservation

- 1** Select the StorageAgent node in the VEA tree view and click **Actions > Clear SCSI reservation**.

Alternatively, you may right-click the Storage Agent node and select **Clear SCSI reservation** on the context menu.
- 2** Review the information message in the window that appears and click **Yes** to complete the operation.

Connection issues

This section describes computer connection issues.

When attempting to connect to a remote computer, you are denied access or get an error message

This topic describes situations where you cannot connect to a remote computer.

Most common situation

In SFW, when you attempt to connect to a remote computer but are denied access, your situation probably falls under one of two conditions.

The conditions are as follows:

- You are logged on to a local server that is not in any domain.
- You are logged on to a domain, but that domain does not contain the remote server. Also your domain does not have a trust relationship with the domain that the remote server is in.

The workaround in this situation is to create an account on the local computer that has the name and password that are the same as the administrator name and password of the remote computer and then use that account to log on to the local computer.

Client-only installation

Another situation where you may get an error message is when you have just done a client-only (VEA) installation for your storage management software. When you start VEA and attempt to connect to a remote server that has Disk Management, Storage Foundation assumes that its client connects first to a remote server running SFW or Volume Manager for Windows before connecting to a system running Disk Management.

Once you connect to a server with SFW or Volume Manager for Windows, you are then able to connect successfully to a remote system running Disk Management.

Note: Windows Disk Management is the disk and volume management program that comes with the Windows operating system. Storage Foundation can connect to a remote server with Disk Management because the Disk Management software is jointly developed by Veritas and Microsoft.

Error message: "The connection to RemoteComputer has terminated. RemoteComputer will be removed from view."

The remote computer that you were connected to has been disconnected from your console. Most often there is a problem with the network connection and the transmission timed out. This can also occur if the remote machine was restarted or the Storage Foundation service on the remote machine was stopped. You can also get the message "Connection Failed" instead.

Check the following:

- Make sure that the remote machine is turned on and available to the network and that the service is started.
- Reconnect to the remote computer.

Issues related to boot or restart

This section contains information about issues that are related to booting or restarting.

See ["During restart, a message may appear about a "Corrupt drive" and suggest that you run autocheck"](#) on page 554.

See ["Error that the boot device is inaccessible, bugcheck 7B"](#) on page 554.

See ["Error message "vxboot- failed to auto-import disk group repltest_dg. all volumes of the disk group are not available.""](#) on page 554.

See ["Cannot extend a dynamic boot or system volume"](#) on page 548.

During restart, a message may appear about a "Corrupt drive" and suggest that you run autocheck

Let autocheck run, but do not worry about the message. Autocheck finishes and the restart is complete. Depending on the size of the system, this may take quite a while.

Error that the boot device is inaccessible, bugcheck 7B

There are two known situations that can result in this message:

After a Dynamic Disk Group with a Boot Device Has Been Renamed

A dynamic disk group that contains a boot volume should not be renamed when the dynamic disk group is imported. Normally, SFW does not let you deport a dynamic disk group that contains the current boot volume. However, if you move a disk to another system or boot from another boot volume in a different dynamic disk group on the same system, you are then able to deport the dynamic disk group that contains the boot volume.

Booting from a boot volume that is contained in a renamed dynamic disk group may cause the error message "Bugcheck 7B, Inaccessible Boot Device."

When Attempting to Boot from a Stale or Damaged Boot Plex

If you get the error message "Bugcheck 7B, Inaccessible Boot Device" and the boot volume is mirrored, try booting off one of the other mirrors. You are not allowed to boot from a stale mirror. A stale mirror is one that was detached because it is missing at some point or because it had errors.

Error message "vxboot- failed to auto-import disk group repltest_dg. all volumes of the disk group are not available."

The situation that causes the error message is that disks from a dynamic disk group in SFW have been removed without being deported, deleted, or converted to basic disks. The next time the computer is restarted, SFW expects to find these disks so that the disk group can be autoimported and when the disks are not there, the error message is generated.

Even when you have uninstalled your storage management software and reinstalled it again, information about the disk group is retained in the registry because it is needed in upgrading from one version of the program to the next.

If you do not plan to use the disk group again, you need to remove the record of that disk group from the registry to avoid getting this message again on restart.

To remove the record of that disk group from the registry:

- 1 Run the command **Regedit** to bring up the Registry Editor window.
- 2 Locate the reference to the disk group in the registry. It is under the following key:
 HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\vxio\Boot
 Info\Disk Group 1, Disk Group 2, ...
- 3 The disk groups are not listed by the actual name, only by a general name of Disk Group 1, Disk Group 2, and so on. To determine which disk group to delete, click the general name to see a further listing that includes the actual name of the disk group.
- 4 Once you have determined which disk group to delete, highlight its general name (such as Disk Group 2) and hit the **Delete** key.

Cluster issues

This section contains information about cluster issues.

See [“Unable to bring a cluster disk group online when a minority of the disks in the disk group are available”](#) on page 555.

See [“Bring a cluster online that has a minority of the disks in the cluster”](#) on page 556.

Unable to bring a cluster disk group online when a minority of the disks in the disk group are available

Safeguards are put into effect so that normally you cannot bring a cluster disk group online on a node when only a minority of disks in the disk group is available. However, in certain situations, you may want to force the import of the disk group or find some other way to bring the disk group online. One example is a campus cluster situation with a two-node cluster, where each node of the cluster is located in a separate physical location and each node has an even number of disks, with the disks evenly distributed between the two sites. In such a situation, if one site goes down, the remaining site cannot access the cluster data because it cannot reserve a majority of the disks in the disk groups.

With Microsoft Failover Clustering and other cluster types, there are methods to allow a site with a minority of the disks to bring disk groups online, if necessary, until the site with failed equipment can be repaired. The following section describes these methods.

Bring a cluster online that has a minority of the disks in the cluster

When bringing a cluster disk group online with a minority of cluster disks, make sure that a majority of the disk group disks are not online on any other cluster node before (and after) onlining the disk group. If a majority of disk group disks are online on another node, data corruption can occur.

Considerations to bring a cluster online:

- If you have a cluster, use the following `vxclus` command for each disk group on your cluster node:

```
vxclus enable -g<DynamicDiskGroupName>
```

You are asked to confirm the use of this command.

If necessary, start the cluster service (`clusvc`) if the cluster service has stopped because of a dynamic quorum resource failure.

Then using the Windows Server Failover Cluster Manager, bring the cluster disk groups online.

More information about the `vxclus` utility is available.

See “[vxclus](#)” on page 621.

- If you have any other cluster type, enter the following command at the command line to receive further information on how to force an import of a cluster disk group when the cluster disk group does not have a majority of disks available.

```
vxdbg -g<DynamicDiskGroupName> -s import
```

A message comes up describing the command for forcing the import and giving cautions about the use of the command.

Enter the command to force the import as follows:

```
vxdbg -g<DynamicDiskGroupName> -s -f import
```

You are asked to confirm the use of this command.

Issues while using SCVMM in an SFW environment

Note: For System Center Virtual Machine Manager (SCVMM) to support operations on volume and disk group resources, the Virtualization Server should have all the latest windows updates installed.

- **Issue while creating virtual machines using SCVMM in an SFW environment:**

This issue is observed if you use System Center Virtual Machine Manager (SCVMM) 2008 R2 or 2008 R2 SP1 to create virtual machines on cluster hosts in SFW.

While creating virtual machines on cluster hosts in an SFW environment, you see the following error message on selecting the virtual machine location path:

```
Cannot create a highly available virtual machine because Virtual Machine Management could not locate or access the volume path
```

Before you use System Center Virtual Machine Manager (SCVMM) 2008 R2 or 2008 R2 SP1 to create virtual machines on cluster hosts in SFW, you must install the following four Microsoft hotfixes on all the cluster hosts to resolve this issue:

KB968936, KB970520, KB971244, and KB971403.

The hotfixes are applied for Windows Server 2008 SP1 and Windows Server 2008 SP2 systems and can be found on the Microsoft website at

<http://support.microsoft.com/kb/2397711/en-us>.

Note: This issue is observed on the Windows Server 2008 SP2 (x64) system.

■ **Issue while adding Hyper-V cluster hosts to SCVMM 2012:**

This issue is observed if you add Hyper-V cluster hosts, which include SFW disk resource type, to the System Center Virtual Machine Manager (SCVMM) 2012.

While adding Hyper-V cluster hosts in an SFW environment, you may see the following error message:

```
Warning (13926)
```

```
Host cluster <Clustername> was not fully refreshed because not all of the nodes could be contacted. Highly available storage and virtual network information reported for this cluster might be inaccurate.
```

To resolve this issue, you must add System Center 2012 Update Rollup 2 patch and install the following Microsoft hotfixes on the host on which SCVMM server and SCVMM Admin Console are configured. You must install these hotfixes before you add cluster hosts to SCVMM 2012:

KB2723538, KB2724539, KB2725034

Information on the hotfixes is provided in the following article on the Microsoft support site:

<http://support.microsoft.com/kb/2706783>.

Dynamic Multi-Pathing issues

This section contains information about the Dynamic Multi-Pathing (DMP) issues.

See [“When a Dynamic Multi-Pathing license is no longer valid”](#) on page 558.

See [“When an unknown disk group appears in a Dynamic Multi-Pathing environment”](#) on page 558.

Note: See the *Veritas InfoScale Release Notes* for documentation of additional issues relating to Dynamic Multi-Pathing. Also, refer to the Hardware Compatibility List on the Veritas Technical Support web site for the list of equipment that has been tested to work with Dynamic Multi-Pathing.

When a Dynamic Multi-Pathing license is no longer valid

If you have a demo license and have DMP DSMs installed on an array and your demo license expires or you replace the demo license key with a permanent license key that does not include Dynamic Multi-Pathing, the program does not let you make changes to your settings. However, all of your current settings remain in effect. You need to purchase a license that enables Dynamic Multi-Pathing to make any additional changes to your settings.

When an unknown disk group appears in a Dynamic Multi-Pathing environment

When an unknown disk group appears in a Dynamic Multi-Pathing environment, it is likely that there is a problem with a path to a storage array or that a storage array was not placed under the control of Dynamic Multi-Pathing. Make sure that all paths are healthy and perform a rescan of all the disks and re-check the status of the disk groups.

vxsnap issues

This section contains a general issue with `vxsnap` and a listing of `vxsnap` error codes with suggested actions to take for each error code.

A description of the `vxsnap` CLI command is available.

See [“vxsnap”](#) on page 710.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxsnap` command.

General issue

When performing a roll-forward recovery to the point of failure with at least 1,008 log files, you may encounter an Exchange initialization error. If you encounter this issue, rename the checkpoint file back to the original name and remount the database stores using Exchange System Manager.

In some cases, updated VSS components are not displayed in VEA console

It is observed that when adding or removing the VSS components or when connecting to the VEA console for the first time, the updated VSS components are not displayed on the VEA console. To view the updated VSS components, you must manually refresh the VEA. To manually refresh the console, run either the `vxsnap refresh` command or perform a Refresh operation through the VEA console.

In VEA GUI, SQL Server database does not show volumes on which they are mounted

This issue is specific to SQL Server 2012 or later. To resolve this, ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each node.

vxsnap command error codes

Occasionally, you may encounter a `vxsnap` error code. For example, the following command can result in the error code 8004230e, if there are conditions that cause the `vxsnap` operation to fail.

```
C:\Documents and Settings\Administrator > vxsnap create source=A:
```

```
Failed to complete the operation...
```

```
8004230e:
```

Error code 8004230e is probably the most common `vxsnap` error code and can occur for several different reasons. In the list below, error code 8004230e and other codes are given, along with suggested actions to take to overcome the error condition.

Error code: 8004230e

VSS snapshot is not supported for the volume(s).

- Possible Reason: Basic partition on a basic disk was used.
 Suggested Action: Create a new dynamic disk group or add the disk to an existing dynamic disk group of SFW.

- Possible Reason: Dynamic disk group version has a version earlier than 4.0 version, when vxsnap was introduced.
 Suggested Action: Upgrade the dynamic disk group to 6.0 or later versions.
- Possible Reason: Snapshot operation is not supported.
 Suggested Action: Install InfoScale Storage. This operation is not supported with InfoScale Foundation.

Error code: 80042316

VSS snapshot already in progress.

Possible reasons and suggested actions:

- Possible Reason: Another VSS job is running.
 Suggested Action: VSS supports only one job at a time. Wait for the other job to finish and then retry the command.
- Possible Reason: A VSS job was abruptly terminated, leaving the VSS Service in an intermediate state.
 Suggested Action: Wait for at least 3 minutes for VSS to reset the state, or, if you are sure there is not another job running, restart the Volume Shadow Copy Service.

Error code: 80042301

The VSS service timed out.

Suggested Action: Wait for at least 3 minutes, then retry.

Suggested Action: When using the `vxsnap` command to snapshot multiple volumes under stress conditions, close the VEA GUI and retry the `vxsnap` command.

Error code: e5090008

The path cannot be used for creating a drive path.

Suggested Action: Make sure that the folder that is used for the drive path is an empty NTFS folder on the local computer.

Other issues

This section contains information about other issues.

On Windows Server 2012 R2, live migration fails if VM VHD is hosted on an SFW volume mounted as a folder mount

On Windows Server 2012 R2 operating systems, an issue occurs while performing a live migration of a Hyper-V virtual machine. If you had mounted the SFW volume as a folder mount and a resource for this volume was used in Hyper-V virtual machine, then the live migration of the virtual machine fails. This happens because, during the live migration, Microsoft's Hyper-V Virtual Machine Management service (VMMS) creates resource mount folder on the target node even before the resource comes online due to which the resource faults and live migration fails.

As a workaround, you need to create a new registry key on all the cluster nodes before performing a live migration. For all the nodes in the cluster, create a DWORD registry key named "VMMSFolderDelete" at the following location and set its value to "1":

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VxSvc\CurrentVersion\VolumeManager
```

On Windows Server 2012, disk group deletion fails if ReFS volume is marked as read-only

On Windows Server 2012 operating system, an issue occurs while deleting a disk group with ReFS volume. If the volume is marked as read-only, SFW is unable to lock the ReFS volume and the volume deletion fails. This in turn causes the disk group deletion to fail.

As a workaround, you can do one of the following:

- Modify the Read only attribute as follows:
 1. In the **VEA GUI** toolbar, right-click the ReFS volume, and then select **Properties** from the context menu.
 2. In the **Properties** window, uncheck the **Read only** attribute and click **OK** to save the change.
- Delete the read-only volume and then delete the corresponding disk group.

Note: This is an operating system issue and not an SFW issue.

On Windows Server 2012, ReFS volume deletion from VEA GUI fails if Symantec Endpoint Protection (SEP) is installed.

On Windows Server 2012 operating systems, an issue occurs while deleting an ReFS volume through the VEA GUI. If SEP is installed, and you try to delete an ReFS volume from the VEA GUI, the volume deletion fails. Due to the ReFS behavior

on the Windows Server 2012 operating system, SFW is unable to lock the ReFS volume and the volume deletion fails.

As a workaround, delete the ReFS volume using the CLI.

An option is grayed out

When an option is grayed out in a menu, the task cannot be performed on the object at this time. Certain operations are valid only for certain types of objects. If there is a task currently running on that object, wait until it has finished and try again. Otherwise, the operation may not be appropriate at this time.

Disk view on a mirrored volume does not display the DCO volume

The Disk View for a selected mirrored volume does not show the DCO volume if that DCO volume was moved to a disk that does not contain any of the other plexes of the mirrored volume. A DCO volume can be moved manually by the user through a subdisk move or it can be moved automatically with the hot relocation function when the subdisk has an I/O error.

If you select the Disk View for all disks (by selecting the Disks folder and clicking the **Disk View** tab in the right pane or by selecting **Disk View** from the Disks folder context menu), you can see the DCO volume. However, finding the DCO volume on all disks may not be practical if you have a large number of disks.

CVM issues

This section helps troubleshoot common Cluster Volume Manager (CVM) problems.

After a storage disconnect, unable to bring volume resources online on the CVM cluster nodes

In a CVM cluster, after storage disconnect occurs and all the nodes re-join the cluster, you may see “Unknown DG” error on all the nodes and disk groups appear in “Deported” state on the Master node. Because of this, you cannot bring the volume resources online on the cluster nodes.

As a workaround, first bring any volume resource online on the Master node. This imports the required disk groups on all the other nodes. Then, bring the volume resources online on the remaining nodes of the CVM cluster.

Error may occur while unconfiguring a node from CVM cluster

In some cases, an error may occur while unconfiguring or removing a node from a CVM cluster using either the SFW Configuration Wizard for Microsoft Failover Cluster or the CLI-based `ffconfigcli.exe` utility.

One of the following error messages may appear:

```
Warning V-194-311-256 Errors were encountered while removing the node
<node name> from the CVM cluster. You may need to clean up the SFW
configuration on the node.
```

```
Error V-194-311-5 An Errors was encountered while stopping the
required service on one or more systems.
```

As a workaround, restart the faulty node, and then unconfigure it using the wizard or the CLI utility.

Shutdown of all the nodes except one causes CVM to hang

In a CVM cluster, if all nodes are shut down except one, then CVM hangs. As all the shutdown nodes come up and try to re-join the cluster, the nodes hang in “Joining” state causing the entire cluster to hang and become non-functional. This affects the CVM functionality on the cluster nodes and does not have any effect on any other processes running on the operating system.

As a workaround, shut down the non-responsive nodes and bring them up one by one and ensure that they become part of the running cluster (both CVM and Microsoft failover cluster).

Sometimes, CSDG Deport causes Master node to hang due to IRP getting stuck in QLogic driver

In some cases, when you perform the disk group deport operation for a cluster-shared disk group (CSDG), it may cause the Master node in the CVM cluster to hang. This happens if I/O request packet (IRP) gets stuck in the QLogic driver. This affects CVM functionality on the cluster nodes and does not have any effect on any other processes running on the operating system.

Veritas needs QLogic’s assistance to resolve this issue.

As a workaround, update the QLogic driver.

Unknown disk groups seen on nodes after splitting a cluster-shared disk group into cluster disk groups from Slave node

When you split a cluster-shared disk group from a Slave node and choose the target disk groups to be cluster disk groups, then the split operation is successful on the Slave node. However, on the other nodes, unknown disk groups are seen because the disk group information is not communicated to those nodes properly.

As a workaround, perform a re-scan on the Master node so that all the nodes have updated information about the disk groups.

In some cases, missing disks are seen on target Secondary dynamic disk groups after splitting a cluster-shared disk group from Slave node

When you split a cluster-shared disk group from a Slave node and choose the target disk groups to be Secondary dynamic disk groups, then the split operation is successful on the Slave node. However, on the other nodes, missing disks may be seen in the disk groups if the disk group information is not communicated to those nodes properly.

As a workaround, run the `vxassist refresh` command on the Master node and the Slave nodes where missing disks are seen. This updates the disk group information on all the nodes and resolve the missing disk issue.

Cannot stop `vxSVC` if SFW resources are online on the node

Using the Service Control Manager or CLI, if you try to stop the Veritas Enterprise Administrator Service (`vxSVC`) while the SFW resources are online on the node, then the service does not stop.

To successfully stop the `vxSVC` service, bring the resources offline, stop the cluster service (`clussvc`), and then stop the `vxSVC` service.

Cluster-shared volume fails to come online on Slave if a stale CSDG of the same name is present on it

From the Master node, if you delete a cluster-shared disk group (CSDG) on a Slave node where the CVM cluster is stopped, then the change is not updated on the Slave node. Now, if you create a new CSDG with the same name as the stale CSDG on Slave, then the cluster-shared volume on Slave fails to come online.

As a workaround, restart the cluster service (`clusssvc`) and the Veritas Enterprise Administrator Service (`vxsvc`) on the node with the stale CSDG by running the following commands in sequence:

```

1 net stop clusssvc
2 net stop vxsvc
3 net start vxsvc
4 net start clusssvc

```

CVM does not start if all cluster nodes are shut down and then any of the nodes are not restarted

If all the nodes of a CVM cluster are shut down simultaneously and then restarted, but one or more nodes do not come up, then CVM does not start.

As a workaround, bring up the shutdown nodes. If the nodes cannot be brought up, then remove them first from the CVM cluster by running the SFW Configuration Wizard for Microsoft Failover Cluster and then from the Microsoft failover cluster.

Incorrect errors shown while creating a CSDG if Volume Manager Shared Volume is not registered

While creating a cluster-shared disk group (CSDG), you may see incorrect errors similar to the following if the Volume Manager Shared Volume resource has not been registered:

V-76-58645-20635: Failed to start CVM cluster on the node

And the node state information may display the following incorrect information:

State: Out of cluster

Reason: network split-brain detected

As a workaround, manually register the Volume Manager Shared Volume resource type. For information about registering the resource manually, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

After splitting or joining disk group having mirrored volume with DRL, VEA GUI shows incorrect volume file system if volumes move to another disk group

This is related to the dynamic disk group that has mirrored volumes with dirty region logging (DRL). When you split or join a dynamic disk group and its volumes are moved to another disk group during the operation, then the VEA GUI incorrectly

shows the volume file system as RAW after the operation is complete. This does not have any effect on I/O operations or product functionality.

As a workaround, perform a Refresh operation so that VEA GUI displays correct file system information.

Enclosure-level storage migration fails, but adds disks if a cluster-shared volume is offline

This is related to migrating data from one enclosure to another for disks with cluster-shared disk groups (CSDGs). If a cluster-shared volume is offline in a disk group, then the whole disk group is not migrated. Migration is partially successful when a disk group with all volumes online is migrated while another disk group with one or more volumes offline is not. Moreover, the operation adds disks to the target enclosure.

As a workaround, manually remove added disks of the target enclosure from the disk group, bring all the cluster-shared volumes online, and then retry the operation.

Volume Manager Shared Volume resource fails to come online or cannot be deleted from Failover Cluster Manager

You have brought a Volume Manager Shared Volume resource offline on a Slave node using Failover Cluster Manager snap-in. Then, you have brought it online on the Master node using VEA GUI. And now, if try to bring it online on the Slave node using Failover Cluster Manager, the resource remains in “Pending” state for a few minutes and then fails. This happens because the volume is already online on the Master node. The following error message appears:

Volume is already online on another node. Offline the volume from the other node and then try again.

Also, if you try to delete the resource during this and Failover Cluster Manager was restarting the resource, then it cannot be deleted. The following error message appears:

Failed to delete <VolumeName>

As a workaround, bring the volume offline on the node where it was brought online through VEA GUI, and then bring it online on the Slave node using Failover Cluster Manager. To delete the resource using Failover Cluster Manager, wait for around 10 minutes and then try the operation.

Sometimes, source cluster-shared volumes are missing after joining two cluster-shared disk groups

In some cases, when you join a cluster-shared disk group (CSDG) with another CSDG, then the cluster-shared volumes of the source disk group are missing in the target disk group.

As a workaround, do one of the following as appropriate:

- If you've already joined two cluster-shared disk groups, then change the internal name of the missing volumes to make them active.
- If you have not performed the join operation, then deport the source cluster-shared disk group and import it as a dynamic disk group, and then join it to the target cluster-shared disk group.

If private CVM links are removed, then nodes may remain out of cluster after network reconnect

In some cases, if the private CVM links across nodes in the cluster gets removed, then after a network reconnect, the nodes may remain out of cluster. The following error is given in the Event Viewer:

```
c00000fb
```

As a workaround, either restart only the affected nodes or restart the CVM stack by running the following commands in sequence:

```
1 vxclustadm stopnode -f
2 net stop llt /y
3 net start vcscomm
4 vxclustadm startnode
```

Format dialog box appears after storage disconnect

In case of storage disconnect when the cluster-shared volume resource faults, you may see the Format dialog box for the volume. The dialog box appears when the storage gets reconnected and the volume online is in the process of coming online. This does not have any effect on your data or product functionality and can be safely ignored.

Volume Manager Shared Volume resources fail to come online on failover nodes if VxSVC is stopped before stopping clussvc

It is not recommended to stop the Veritas Enterprise Administrator Service (vxsvc). However, if required, then you must stop the clussvc service first, and then stop the vxsvc service. If you stop vxsvc before stopping clussvc on the node, then the Volume Manager Shared Volume resources may fail to come online on the failover nodes.

As a workaround, start the vxsvc service on the node where it was stopped.

One or more nodes have invalid configuration or are not running or reachable

This is regarding various CVM cluster configuration-related issues where one or more cluster nodes have invalid or are not running or reachable.

This section provides information on the following two error messages:

```
Error V-194-311-701: One or more systems in the cluster have an
invalid configuration. For details, select a system and view its
status details or see the log file. Resolve the configuration issues
on such systems or remove them from Microsoft failover cluster, and
then retry the operation.
```

This error occurs for one or more of the following reasons:

- WMI issues
- Connectivity/authentication issues
- One or more nodes are not running or reachable
- Firewall issues

For the reasons listed above, the SFW Configuration Wizard for Microsoft Failover Cluster or the ffconfigcli utility do not let you perform any further operations. To resolve these issues, do one or more of the following as required:

- 1 **Recommended Solution:** Try to resolve the connectivity issues and bring up the affected node. Then, re-launch the wizard and perform the desired operations.
- 2 If a single node cannot be brought up, then do the following:
 - Remove the node from the Microsoft failover cluster.
 - Remove the node from the CVM cluster by running the ffconfigcli `-removenode <NodeName>` command.

- If you want to use this node again, then you first need to clean up the configuration by running the `ffconfigcli -removenode <NodeName>` command on the node.
- 3** If multiple nodes cannot be brought up, then do the following:
- Remove the nodes from the Microsoft failover cluster.
 - Bring all the running applications offline.
 - Unconfigure the CVM cluster using the wizard or by running the `ffconfigcli -unconfigure` command.
 - Re-configure the CVM cluster with the healthy nodes using the wizard or by running the `ffconfigcli -autoconfigure` command.
 - Bring all the applications online.
 - If you want to use these nodes again, then you first need to clean up the configuration by running the `ffconfigcli -removenode <NodeName>` command on each node.

Warning V-194-311-256: Errors were encountered while removing the node CVMCLUSNODE4 from the CVM cluster. You may need to clean up the SFW configuration on the node. This error occurs if the node that you were trying to remove is not running or reachable.

This error occurs if the node that you try to remove is not running or reachable.

To resolve this issue, once the node is up, remove the faulty node by re-launching the wizard or by running the `ffconfigcli -removenode <NodeName>` command on the node.

After node crash or network disconnect, volume resources failover to other node but the drive letters are left behind mounted on the failing node even after it joins cluster successfully

In CVM, after a node crashes or network gets disconnected, the cluster-shared volume resources are failed over to another healthy node. If you change the drive letters for volumes on that node and perform a failover to the original node when it is up, then the cluster-shared volume resources fail to come online. This happens because the affected node retains the older drive letters even if the mount points were changed on another node.

As a workaround, manually remove the older drive letters from the affected node.

Shutdown of Master node in a CVM cluster makes the Slave nodes to hang in “Joining” state while joining to new Master

When the Master node of a CVM cluster is shut down, then one of the Slave nodes becomes the new Master. Because of this, the remaining nodes try to re-join as Slave to the new Master node. If there is an issue with GAB message handling between the new Master and Slave nodes, then the Slave nodes hang in “Joining” state.

When you use the `vxclustadm [-v] nodestate` command for node information, it displays "Vxconfigd in join" in the output.

As a workaround, restart the VxSVC service on the hung Slave nodes.

CVM stops if Microsoft Failover Clustering and CVM cluster networks are not in sync because of multiple, independent network failures or disconnect

This issue exists because Microsoft Failover Clustering and GAB form two different cluster networks for the CVM cluster nodes. If there is a multiple, independent network failure or disconnect on different nodes, then there will not be a common CVM link to be shared by the nodes across cluster networks.

Because of this, the cluster networks go out of sync and CVM stops working.

As a workaround, do the following as required:

- Resolve the network connectivity issues.
- If new network links have to be set up, then do as follows: Configure the new network links, and then unconfigure and reconfigure the CVM cluster to reflect the network link changes.
Please note that this requires an application downtime.

Restarting CVM

Sometimes, you may need to restart Cluster Volume Manager (CVM) to reflect certain changes. Using the command line interface (CLI), run the following commands sequentially to stop and then start a CVM cluster:

Commands to stop CVM:

- `vxclustadm stopnode -f`
- `net stop llnt /y`

Commands to start CVM:

- `net start vcscomm`

- `vxclustadm startnode`

Administering CVM using the CLI

Apart from the VEA GUI, you can perform administrative tasks of CVM using the command-line interface (CLI). Among the available commands, the most useful is the `vxclustadm` command, which can be used to start or stop a node in a CVM cluster, view the state of a node, or view the node map.

Note: It is recommended that you use the `vxclustadm` command only when it is most necessary.

For information on using `vxclustadm` and other commands to administer CVM:

See [“SFW commands available from the command line”](#) on page 574.

Tuning the VDS software provider logging

New registry keys are added to tune the Virtual Disk Service software provider logging. The VDS software provider can be tuned using the following details:

- Location: `%vmpath%/log`
- File names: `vxvds.log` and `vxvdsdyn.log`

The following two registry keys are used to tune the VDS software provider logging:

- **MAXSIZE** gives the maximum size of an individual log file in KB units. A backup log file is created when the log file exceeds the default size. The default value of **MAXSIZE** is 16384 KB. However, you can customize the default value.
 Run the command `Regedit` to open the registry editor.
 Locate the **MAXSIZE** of the VDS software provider in the registry under the following key:
`SOFTWARE\Veritas\VxSvc\CurrentVersion\Tracing\vds\MaxSize.`
- **MAXFILES** gives the maximum number of log files that can be present at a time. The old log files are deleted once they exceed the limit of maximum number of files. The default value of **MAXFILES** is 5. However, you can customize the default value.
 Run the command `Regedit` to open the registry editor.
 Locate the **MAXFILES** of the VDS software provider in the registry under the following key:

```
SOFTWARE\Veritas\VxSvc\CurrentVersion\Tracing\vds\MaxFiles
```

For example, the default log file is `vxvds.log`. If the created log file exceeds **MAXSIZE**, it is copied to `vxvds1.log`, and `vxvds.log` is erased. New logs are appended to the `vxvds.log` file.

Command line interface

This appendix includes the following topics:

- [Overview of the command line interface](#)
- [vxclustadm](#)
- [vxvol](#)
- [vxdg](#)
- [vxclus](#)
- [vxdisk](#)
- [vxassist](#)
- [vxassist \(Windows-specific\)](#)
- [vxevac](#)
- [vxsd](#)
- [vxstat](#)
- [vxtask](#)
- [vxedit](#)
- [vxunreloc](#)
- [vxdmpadm](#)
- [vxcbr](#)
- [vxsnap](#)
- [vxfsync](#)

- [vxscrub](#)
- [vxverify](#)
- [vxprint](#)
- [sfcache](#)
- [Tuning SFW](#)

Overview of the command line interface

This section is an overview of the command line interface.

SFW commands available from the command line

This section describes the command line interface (CLI) commands available with Storage Foundation. The key SFW commands can be executed from the command line as well as from the GUI.

The utilities are available in the product installation directory, typically `C:\Program Files\Veritas\Veritas Volume Manager`. You can run the commands from any directory in the command prompt window.

Note: The CLI commands run only on the server. They do not run on the Storage Foundation client.

Note: Storage Foundation Manager does not support the CLI commands.

Note: If User Access Control (UAC) is enabled, then run the CLI commands in the “Run as administrator” mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

The available commands are the following:

- [vxvol](#)
Displays volume information, repairs a RAID-5 or mirrored volume, turns Fast Resynchronization on and off, and reactivates a volume. See “[vxvol](#)” on page 581.
- [vxdg](#)
Lets you create a dynamic disk group and perform other functions that are related to dynamic disk groups. See “[vxdg](#)” on page 593.

- **vxclus**
 Enables and disables the capability of forcing the import of a disk group on a cluster that is controlled by Microsoft Failover Clustering. The command also has keywords that allow the cluster disk group to reside on the same bus as the cluster node's system and boot disk. See "[vxclus](#)" on page 621.
- **vxclustadm**
 Enables you to start and stop a node in a Cluster Volume Manager (CVM) cluster, view the state of a node, and view the mapping and role (Master or Slave) information of the nodes. See "[vxclustadm](#)" on page 580.
- **vxdisk**
 Provides information on dynamic disks, lets you reactivate a dynamic disk, and enables track alignment for the disk. See "[vxdisk](#)" on page 624.
- **vxassist**
 Lets you create a dynamic volume and perform other operations that are related to dynamic volumes such as shrinking the size of the volume. See "[vxassist](#)" on page 638.
- **vxevac**
 Evacuates (moves) subdisks from the disk that is specified to a specified destination disk. See "[vxevac](#)" on page 678.
- **vxunreloc**
 Reverses the process of hot relocation. See "[vxunreloc](#)" on page 689.
- **vxsd**
 Provides commands for moving and splitting a subdisk and joining two subdisks. See "[vxsd](#)" on page 679.
- **vxstat**
 Provides commands for displaying statistics for disks, volumes, and subdisks. See "[vxstat](#)" on page 681.
- **vxtask**
 Provides commands for displaying percentage completion information, and for pausing and resuming the tasks. See "[vxtask](#)" on page 684.
- **vxedit**
 Provides commands to rename or add comments to a dynamic disk group, a disk, a volume, a plex, and a subdisk. See "[vxedit](#)" on page 687.
- **vxdmppadm**
 Provides commands for Dynamic Multi-Pathing (DMP DSMs). See "[vxdmppadm](#)" on page 690.
- **vxabr**

Provides commands for backing up and restore the Storage Foundation for Windows configuration information. See “[vxabr](#)” on page 708.

- **vxsnap**
Performs multiple snapshots simultaneously. The utility integrates with the Windows Volume Shadow Copy Service (VSS) as a VSS Requester. This allows for the simultaneous snapshot of all volumes that are associated with an Exchange Server database or an SQL Server database. See “[vxsnap](#)” on page 710.
- **vxfsync**
Resynchronize a single file in a snapshotted volume back to the original volume. It is designed to restore one or more corrupt files in a database application that has volumes with multiple database files. See “[vxfsync](#)” on page 743.
- **vxscrub**
Removes unneeded entries in the registry that are a result of adding and deleting volumes. The command is designed to shrink the size of the registry by purging the system hive. See “[vxscrub](#)” on page 743.
- **vxverify**
Determines if there are any differences between plexes of mirrored volumes. See “[vxverify](#)” on page 746.
- **vxprint**
Displays information about SFW and Volume Replicator objects. See “[vxprint](#)” on page 747.

Command line log

SFW provides support for a log that captures commands that are issued through the CLI. The command line log also captures the system response to each command. The log file, `vxcli.log`, is typically located at `C:\Program`

`Files\Veritas\Veritas Volume Manager\logs`.

Each entry in the log is prefixed with the date and time of the command or system response, followed by the process identification and thread identification. Each CLI command is a separate entry in the log and each is identified by "----- Command : ". Each system response is also a separate entry in the log and follows the CLI command.

Note: Help messages are not recorded in the command line log.

As an example of a log's content, the following was recorded as a result of entering `vx dg list`

```

03/28/06 22:59:37 [5376,2024] ----- Command : vxdglist
03/28/06 22:59:38 [5376,2024] Disk groups are....
03/28/06 22:59:38 [5376,2024] Name : BasicGroup
03/28/06 22:59:38 [5376,2024] Disk Group Type : Basic
03/28/06 22:59:38 [5376,2024] DiskGroup ID : No ID
03/28/06 22:59:38 [5376,2024] Status : Imported
03/28/06 23:00:08 [4212,5760] ----- Command : vxdisklist
03/28/06 23:00:09 [4212,5760] Name MediaName DiskgroupDiskStyle Size (MB)
FreeSpace (MB) Status
03/28/06 23:00:09 [4212,5760] Harddisk0 BasicGroup MBR117239 67240
Uninitialized
03/28/06 23:00:09 [4212,5760] Harddisk1 BasicGroup MBR1019 917
Uninitialized
03/28/06 23:00:09 [4212,5760] Harddisk2 (No Signature)BasicGroup
MBR 10236 10236 Uninitialized
03/28/06 23:00:09 [4212,5760] Harddisk3 BasicGroup MBR1019 917
Uninitialized
03/28/06 23:00:09 [4212,5760] Harddisk4 (No Signature)BasicGroup
MBR 10236 10236 Uninitialized
03/28/06 23:01:41 [1100,4728] ----- Command :
vxdg-g dgl init harddisk1
03/28/06 23:01:47 [1100,4728] Successfully createddynamic
disk group dgl

```

Conventions for command line syntax

This topic describes the typographical conventions that are used for the command line syntax in this CLI section.

The conventions are as follows:

- Any parameter that is optional for the command syntax has square brackets ([]) around it. For example:
`[-b] or [-o keepplex]`
- Required command words and parameters for the command do not have square brackets around them. For example:
`vxvol volinfo or <VolumeName>`
- Command words and parameters that are typed as shown in the command syntax are displayed in the Courier bold font. For example:
`vxvol volinfo or [-b]`
- Parameters that require the user to enter something in their place are displayed in Helvetica Italic font and have angle brackets around them. They are placeholders for information the user enters. For example:

<VolumeName>

If a parameter is enclosed within both angle brackets and square brackets, it is an optional parameter for the command. For example:

[<DiskName>]

- The | character is a separator that allows two or more choices for a given parameter. The user can use any one of the choices for the command. For example:
 <VolumeName | DriveLetter>
 [volume|mirror|log]
- Ellipses (...) after a parameter indicate more items. For example, the parameter <DiskName>... means to enter one or more disk names. The parameter [DriveLetter={A|B|...|Z}] means to enter DriveLetter= and a drive letter in the range of A to Z.
- The parameter !<DiskName> is used with three commands—vxassist break, vxassist remove, and vxassist snapshot. The exclamation mark is needed to make these commands consistent with the Volume Manager for UNIX versions of these commands.
- About Volume Names
 - Assigning a Volume Name
 When you create a volume with vxassist make, the <VolumeName> parameter is a name you assign to the volume.

Note: A volume name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.

The volume name you assign in either the command line or the GUI is the internal name of the volume. SFW makes that internal volume name into the operating system volume label. However, if you later change the volume label through the operating system, the internal SFW volume name does not change.

Note: The name you assign for the <VolumeName> parameter when creating a volume in the CLI is equivalent to the volume name that is specified in creating a volume with the New Volume command in the VEA GUI.

The internal volume name supplies an identifier for the volume that stays the same. The operating system may reassign drive letters. A persistent

name is especially important in scripts. If you want to use the internal volume name in the command line to refer to the volume, you must precede it with its disk group parameter, `-g<DynamicDiskGroupName>`, for example, `-gDG1 Salesvolume`.

If you do not assign a volume name, SFW assigns the name, in the form of `volume1`, `volume2`, etc. The program assigns the internal volume name only to dynamic volumes.

- Other Ways to Designate a New Volume

A basic or dynamic volume can also be indicated by a drive letter, `<DriveLetter>`. If the volume has been mounted at an NTFS folder, the volume name is indicated by the drive path parameter `<DrivePath>`, which is the path to the folder.

- Referring to an Existing Volume

There are four methods for referring to an existing volume in the CLI:

- The full path name of the volume, which has the following syntax:

```
\Device\HarddiskDmVolumes\<DynamicDiskGroupName>\
<VolumeName>
```

For example:

```
\Device\HarddiskDmVolumes\DynDskGrp1\Volume1
```

- With a drive letter, such as `D`:
- With its internal volume name, which requires the disk group option
 For example, `-gDG1 Salesvolume` or `-gDG1 Volume1`
- With a `<DrivePath>` parameter if the volume is mounted on an NTFS folder, such as `D:\Sales`.

Note: In commands that involve existing volumes, the internal volume name is referred to as `<VmName>` in the command line syntax.

You can identify the internal volume name through the `vxvolinfo` command, or you can see the internal volume name in the right pane of the GUI when a volume is selected.

- About Disk Names

A disk can be specified in the command line in different ways:

- With the device name, such as `Harddisk2`

- With a path name, such as `\Device\Harddisk2`
 - With the internal disk name assigned by the SFW program, such as `Disk2`
 - With `p#c#t#l#`, where the #’s correspond to port, channel, target, and LUN
- The internal disk name is a persistent name that remains even when the operating system changes the disk name, and thus it is useful in scripts. It requires a disk group parameter whenever it is used (for example, `-gDG1 Disk2`). You cannot assign this name. You can identify a disk’s internal name by using the `vxdisk diskinfo` command. An internal disk name is assigned only to dynamic disks, not basic disks.
- You can also see the device name and the internal disk name on the right pane of the VEA GUI when a disk or disks are selected.

vxclustadm

<code>startnode</code>	To start the CVM cluster on a node.
<code>stopnode</code>	To stop the CVM cluster on a node.
<code>nodestate</code>	To view the state of a particular node in the CVM cluster.
<code>nidmap</code>	To view node mapping and role (Master or Slave) information of the nodes in the CVM cluster.

vxclustadm startnode

```
vxclustadm startnode
```

Starts CVM on the node where you run this command.

vxclustadm stopnode

```
vxclustadm [-f] stopnode
```

Stops CVM on the node where you run this command.

vxclustadm nodestate

```
vxclustadm [-v] nodestate
```

Where `-v` stands for verbose.

Gives information about the state of a particular node in the CVM cluster.

vxclustadm nidmap

```
vxclustadm nidmap
```

Gives information about the node mapping and role (Master or Slave) of the nodes in the CVM cluster.

vxvol

<code>volinfo</code>	Displays information on a dynamic volume.
<code>repair</code>	Repairs a RAID-5 volume.
<code>set fastresync=on off</code>	Turns Fast Resynchronization on and off.
<code>reactivate</code>	Reactivates the volume (restarts the volume).
<code>rdpol prefer</code>	Sets the volume read policy on a volume with multiple mirrors to designate one specific plex to be used whenever possible for reading from the volume.
<code>rdpol round</code>	Sets the volume read policy on a volume with multiple mirrors so that plexes are read sequentially in a "round-robin" fashion.
<code>growfs</code>	Grow the file system.
<code>access</code>	Sets the volume to read only or read/write access mode.
<code>reclaim</code>	Reclaim storage space from a volume. (Portion of volume must reside on a thin provisioned disk.)
<code>reclaimcancel</code>	Immediately cancels reclaim operation
<code>capacitymonitor</code>	Set threshold on dynamic volume so that when the used disk space on a volume reaches the specified threshold, an alert message is logged.
<code>autogrow</code>	Allows an individual volume to grow (increase capacity) automatically by a specified amount when the used disk space on the volume reaches a certain threshold. The threshold is set by Capacity Monitoring.

The volume name or drive letter follows each keyword. Typing the following sequence for each keyword brings up a description of its syntax:

```
vxvol <Keyword> -?
```

vxvol volinfo

```
vxvol [-v] [-g<DynamicDiskGroupName>] volinfo <VolumeName |  
DriveLetter | VmName | DrivePath>
```

Returns information about the volume.

The following attributes apply:

<code>[-v]</code>	The verbose option displays the size and offset of the subdisks as well as the other information provided by the main command.
<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\Dg1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDg1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

The possible information categories are name, dynamic disk group name, size, volume type, whether it is a boot or system volume, whether it contains the page file (pagefile.sys), volume status, and access mode (read only or read/write access mode). If the volume is a dynamic volume, the command can return additional information on any mirror plexes or DCO volumes that are associated with the volume.

Example

```
vxvol volinfo E:
```

This returns a list with specified information as described above for volume E.

vxvol repair

```
vxvol [-o notrackalign] [-g<DynamicDiskGroupName>] repair <VolumeName  
| DriveLetter | <VmName | DrivePath>
```

Repairs a RAID-5 or mirrored volume if its status is Degraded and there is unallocated space on another dynamic disk. The command repairs the damaged part of the volume by placing a functional replacement of that subdisk on another disk. The following attributes apply:

<code>-o notrackalign</code>	Disable track alignment on the destination disk.
<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\Dg1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDg1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Example

```
vxvol repair J:
```

This command repairs a RAID-5 volume (designated as J:) after a disk that belongs to the RAID set has failed.

vxvol set fastresync=on|off

```
vxvol [-o notrackalign] [-g<DynamicDiskGroupName>] set
fastresync=on|off <VolumeName | DriveLetter |<VmName|DrivePath>
```

Turns FastResync on or off for the specified mirrored volume. The following attributes apply:

<code>-o notrackalign</code>	Disable track alignment on the disk where the log resides. Applicable when <code>fastresync=on</code> .
<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name.

<i><VolumeName></i>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<i><DriveLetter></i>	The drive letter of the volume.
<i><VmName></i>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1.</code>
<i><DrivePath></i>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

FastResync is not supported on RAID-5 volumes. If you have initiated a snapshot operation on a volume, you cannot turn off FastResync for that volume. If you try to do so, the command line interface returns an error message.

When you turn FastResync on, a DCO (disk change object) log volume is created. When you turn off FastResync, the DCO log volume is deleted. You can perform a limited number of operations on a DCO log volume.

The following topics provide more information:

See [“FastResync”](#) on page 253.

See [“Disk Change Object \(DCO\) volume”](#) on page 256.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Examples

```
vxvol set fastresync=on J:
```

This command turns on the FastResync feature on the volume with drive letter J.

```
vxvol set fastresync=on\Device\HarddiskDMVolumes\DynDskGrp1\Volume1
```

This command turns on the FastResync feature for Volume 1, which belongs to DynDskGrp1.

vxvol reactivate

```
vxvol [-g<DynamicDiskGroupName>] reactivate <VolumeName | DriveLetter  
| <VmName | DrivePath>
```

Manually restarts the volume. You need to rescan and then reactivate the disk before attempting to reactivate the volume.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxvol reactivate` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name parameter.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

More information about the CLI version of the Rescan command and the Reactivate (a disk) command is available.

- See “[vxassist](#)” on page 638.
- See “[vxdisk](#)” on page 624.

Example

```
vxvol reactivate J:
```

This command reactivates the volume with drive letter J.

vxvol rdpol prefer

```
vxvol [-g<DynamicDiskGroupName>] rdpol prefer <VolumeName |  
DriveLetter |<VmName|DrivePath><PlexName>
```

Sets the volume read policy on a volume with multiple mirrors to designate a specific plex to be used for reads. This plex is referred to the "preferred plex."

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxvol rdpol prefer` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName##gt##</code>	Needed only if you use the internal volume name parameter.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDr\Volumes\DG1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code><PlexName></code>	Name or GUID of the desired preferred plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command <code>vxvol volinfo</code> for the mirrored volume that contains the plex.

Designating a plex as preferred can improve your system's read performance if the drive used for the target mirror is a high-performance drive or if the target mirror is located on a drive that is physically closer.

Note: When a preferred plex has an error, another plex completes the read operation. This happens automatically, as if there were no preferred plex.

Example

```
vxvol rdpol prefer J: Volume1-01
```

This command designates plex Volume1-01 as the preferred plex on mirrored volume J.

vxvol rdpol round

```
vxvol [-g<DynamicDiskGroupName>] rdpol round <VolumeName | DriveLetter  
| <VmName | DrivePath>
```

Sets the volume read policy on a volume with multiple mirrors to read from each plex sequentially in a "round-robin" fashion. Sequential access causes only one

plex to be accessed at a time. This takes advantage of the drive or controller read-ahead caching. The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name parameter.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Example

```
vxvol rdpol round J:
```

This command sets up the mirrors of volume J to be read in a round-robin fashion.

vxvol growfs

```
vxvol [-g<DynamicDiskGroupName>] growfs <VolumeName | DriveLetter  
| <VmName | DrivePath> [<length>]
```

Grows the file system if a volume is resized but the file system fails to resize correctly. If the optional parameter [`<length>`] is not used, the command tries to grow the file system size to the full volume size.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name parameter.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

<length> Size of the volume in MB. If no size is indicated, the volume is grown to the maximum size.

Note: Any growth is limited to a maximum of 2 TB every time this command is run. The file system usable space is in general somewhat smaller than the raw volume size, since the file system itself occupies space. Only use this command when you see a large difference between the volume size and the file system capacity.

Example

```
vxvol growfs J: 1000
```

This command would grow the file system size of volume J by 1 GB.

```
vxvol growfs J:
```

This command would grow the file system size of volume J to the full volume size.

vxvol access

```
vxvol [-g<DynamicDiskGroupName>] [-f ]
access readonly|readwrite <VolumeName | DriveLetter |<VmName|DrivePath>>
```

Sets the read or write access mode of a volume. The following attributes apply:

<i>-g<DynamicDiskGroupName></i>	Name of dynamic disk group. (Needed only if you use the internal volume name parameter.)
<i>[-f]</i>	The force option forces the volume to the specified access mode.
<i><VolumeName></i>	The path name of the volume, such as \Device\HarddiskDmVolumes\DG1\Volume1.
<i><DriveLetter></i>	The drive letter of the volume.
<i><VmName></i>	Internal name of the volume; requires the use of the <i>-g</i> option, such as <i>-gDG1 Volume1</i> .
<i><DrivePath></i>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Example

```
vxvol access readwrite E:
```


This command sets volume E to read or write access.

vxvol reclaim

```
vxvol [-g<DynamicDiskGroupName>] [-b]  
reclaim<VolumeName|DriveLetter|VmName|DrivePath>
```

Reclaim storage space from a volume. (Portion of volume must reside on a thin provisioned disk.)

The following attributes apply:

<i>-g</i> <DynamicDiskGroupName>	Name of dynamic disk group. (Needed only if you use the internal disk name of the volume.)
<i>-b</i>	Run the command in the background.
<i>VolumeName</i>	The path name of the volume, such as \Device\HarddiskDmVolumes\DG1\Volume1
<i>DriveLetter</i>	Drive letter of the volume.
<i>VmName</i>	The internal disk name of the volume, which requires the <i>-g</i> option (for example, <i>-gDG1Volume2</i>).
<i>DrivePath</i>	Drive path that the volume is mounted on. (Must be an empty folder on a local NTFS volume.)

Note: The operation to reclaim storage from thin provisioned disks is not supported for RAID-5 volumes.

Note: The reclaim storage operation is disabled during a shrink volume operation. The reclaim storage operation is enabled when the shrink volume operation has completed.

Example

```
vxvol -gDG1 reclaim Volume1
```

This command reclaims space from volume with internal name Volume1.

vxvol reclaimcancel

```
vxvol [-g<DynamicDiskGroupName>]  
reclaimcancel<VolumeName|DriveLetter|VmName|DrivePath>
```

Immediately cancels the reclaim operation. Only a portion of the storage space of the volume may have been reclaimed.

The following attributes apply:

<i>-g</i> <DynamicDiskGroupName>	Name of dynamic disk group. (Needed only if you are using the internal disk name of the volume.)
<i>VolumeName</i>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code>
<i>DriveLetter</i>	Drive letter of the volume.
<i>VmName</i>	The internal disk name of the volume, which requires the <code>-g</code> option (for example, <code>-gDG1Volume2</code>).
<i>DrivePath</i>	Drive path that the volume is mounted on. (Must be an empty folder on a local NTFS volume.)

Example

```
vxvol reclaimcancel F:
```

This command immediately cancels the reclaim operation on volume that is associated with drive letter F:.

vxvol capacitymonitor

```
vxvol [-g<DynamicDiskGroupName>]  
capacitymonitor<VolumeName|DriveLetter|VmName|DrivePath>  
enable=<yes|no> [critical=<# of volume capacity>] [warning=<# of  
volume capacity>] [pollinginterval=<# of seconds>] [messageinterval=<#  
of seconds>] [email=<email address>] [-A APPLYPOLICYTOFAILOVER]
```

Set threshold on dynamic volume so that when the used disk space on a volume reaches the specified threshold, an alert message is logged.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal disk name of the volume.)
<code>VolumeName</code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code>
<code>DriveLetter</code>	Drive letter of the volume.
<code>VmName</code>	The internal disk name of the volume, which requires the <code>-g</code> option (for example, <code>-gDG1Volume2</code>).
<code>DrivePath</code>	Drive path that the volume is mounted on. (Must be an empty folder on a local NTFS volume.)
<code>Enable</code>	Enable or disable capacity monitoring.
<code>Critical</code>	Critical threshold. The percentage of volume capacity when SFW reports an error. The default value is 90% of capacity.
<code>Warning</code>	Warning threshold. The percentage of volume capacity when SFW logs a warning alert
<code>Email</code>	Notification email address when volume capacity has reached warning/critical threshold.
<code>PollingInterval</code>	Polling interval for checking capacity. The polling interval is the interval of time, in seconds, at which Storage Foundation for Windows checks volume capacities. The default value is 20 seconds.
<code>MessageInterval</code>	Minimum time between alerts. To keep the event log file from getting full, you can limit how often an alert message is generated after the warning threshold is reached for a volume. The default is 100 seconds between messages
<code>-A APPLYPOLICYTOFAILOVER</code>	Make capacity monitoring policy available to another cluster node after a failover.

vxvol autogrow

```
vxvol [-g<DynamicDiskGroupName>]  
autogrow<VolumeName|DriveLetter|VmName|DrivePath> enable=<yes|no>
```

```
[growby=<length>] [-t tpdiskonly] [maxvolumesize=<#size>]  
[email=<email address>] [<alloc_attributes...>]  
[<diskname|p#c#t#l#>...] [-A APPLYPOLICYTOFAILOVER]
```

Allows an individual volume to grow (increase capacity) automatically by a specified amount when the used disk space on the volume reaches a certain threshold. The threshold is set by Capacity Monitoring.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal disk name of the volume.)
<code>VolumeName</code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DGI\Volume1</code>
<code>DriveLetter</code>	Drive letter of the volume.
<code>VmName</code>	The internal disk name of the volume, which requires the <code>-g</code> option (for example, <code>-gDGI Volume2</code>).
<code>DrivePath</code>	Drive path that the volume is mounted on. (Must be an empty folder on a local NTFS volume.)
<code>Enable</code>	Enable or disable automatic volume growth.
<code>growby</code>	Amount of space that volume is to be expanded each time the growth threshold is reached. Amount of space can be specified in terms of S for Sectors, K for KBs, M for MBs, G for GBs, T for TBs. The default unit of measurement is MBs.
<code>-t tpdiskonly</code>	Specify for only thin provisioned disks.
<code>maxvolumesize</code>	Maximum volume size. Limit volume growth to a maximum volume size. Maximum volume size can be specified in terms of S for Sectors, K for KBs, M for MBs, G for GBs, T for TBs. The default is unrestricted volume growth.
<code>email</code>	Notification email address when volume capacity has reached warning/critical threshold.

<code><alloc_attributes...></code>	Specify allocation attributes for mirror/striped volumes. Attributes are of the form: mirror={diskclass} wantmirror={diskclass} stripe={diskclass} wantstripe={diskclass}
<code>Mirror</code>	Specify the diskclass to use for mirroring.
<code>WantMirror</code>	Specify the diskclass to use for mirroring. This specification is ignored if allocations are not available.
<code>Stripe</code>	Specify that volumes are to be striped across instances of a diskclass.
<code>WantStripe</code>	Specify that volumes are to be striped across instances of a diskclass. This specification is ignored if allocations are not available.
<code>diskclass</code>	Type of disk grouping to be used for mirror/striped. Available types are channel, enclosure, port, or target.
<code>diskname</code>	Name of the disk, e.g., Harddisk2; or internal disk name (requires use of -g option), e.g., -g MyDg Disk2
<code>p#c#t#l#</code>	Corresponds to the Port, Channel, Target, LUN of a disk.
<code>-A APPLYPOLICYTOFAILOVER</code>	Make capacity monitoring policy available to another cluster node after a failover.

vxdg

<code>list</code>	Displays a list of the dynamic disk groups on the computer.
<code>list (CSDG)</code>	Displays a list of the cluster-shared disk groups (CSDGs) on the computer.
<code>dginfo</code>	Displays information about a specified dynamic disk group.
<code>dginfo (CSDG)</code>	Displays information about a specified cluster-shared disk group (CSDG).
<code>init</code>	Creates a dynamic disk group.

<code>init (CSDG)</code>	Creates a cluster-shared disk group (CSDG).
<code>adddisk</code>	Adds a basic disk to a dynamic disk group.
<code>rmdisk</code>	Removes a disk from a dynamic disk group and reverts it back to a basic disk.
<code>import</code>	Imports the specified dynamic disk group on the computer.
<code>import (CSDG)</code>	Imports the specified cluster-shared disk group (CSDG) on the computer.
<code>deport</code>	Deports the specified dynamic disk group on the computer.
<code>destroy</code>	Deletes the specified dynamic disk group on the computer.
<code>protect</code>	Adds private dynamic disk group protection.
<code>release</code>	Removes private dynamic disk group protection.
<code>upgrade</code>	Upgrades the disk group version to the current version (the default) or earlier version of Volume Manager.
<code>repldisk</code>	Replaces the specified disk by moving all the subdisks to a spare disk.
<code>split</code>	Splits the specified dynamic disk group into two dynamic disk groups.
<code>recover</code>	Recovers a dynamic disk group that fails because of a system crash or other problem during a dynamic disk group split operation.
<code>join</code>	Joins two dynamic disk groups into one larger dynamic disk group.
<code>latestart</code>	Imports disk groups after system starts up to allow more time for storage to become available for applications.
<code>reclaim</code>	Reclaim storage space from thin provisioned disks in a dynamic disk group.
<code>reclaimcancel</code>	Immediately cancel reclaim operation
<code>refreshff</code>	Refreshes the disk group state on the system to support the new Deported Read-Only state for fast failover.

A dynamic disk group is identified by `-g<DynamicDiskGroupName>`, such as `DynDskGrp1` or by its dynamic disk group ID (DgID). The DgID is an internal number assigned to the disk group. It can be viewed through the `vxdg list` or `vxdg dginfo` command.

`<DiskName>` or `p#t#l#` (where the #s corresponds to the port, channel, target, and LUN of a disk) identifies a disk that is added or removed.

In early releases of Volume Manager for Windows, using `vxdg adddisk` to add the first basic disk to a dynamic disk group automatically created the first dynamic disk group (known as the primary dynamic disk group). If you then used `vxdg adddisk` to specify adding a disk to a dynamic disk group with a new name, a secondary dynamic disk group was formed. You must use the `vxdg init` command to create a dynamic disk group. The `vxdg adddisk` command now only adds disks to the dynamic disk groups that have already been created.

Note: In Storage Foundation for Windows, the primary dynamic disk group is the dynamic disk group that contains the boot or system disk. It is not necessarily the first dynamic disk group that is created.

Note: In all versions of Volume Manager, occasionally if volumes arrived after commands like `import`, `init`, `adddisk`, and `join` are completed, subsequent commands like associating a drive letter might fail. However, in the Storage Foundation for Windows, these commands wait until the volumes are ready to be used. If the volumes take a very long time to arrive (a rare case), the command may timeout so that the script does not hang. Users can use `-o timeout=<n>` to override the default timeout.

Typing the following sequence for each keyword brings up a description of its syntax:

```
vxdg <keyword> -?
```

vxdg list

```
vxdg list
```

Returns a list of the dynamic disk groups that are on the current system. The list includes the disk groups' names, their types, whether each disk group is imported or deported, and the DgID. The dynamic disk group ID (DgID) is an internal number assigned to the disk group. It can be used in place of the dynamic disk group name.

vxdg list (CSDG)

```
vxdg -csdg list
```

Returns a list of the cluster-shared disk groups (CSDGs) that are on the current system. The list includes the cluster-shared disk groups' names, their types, whether each disk group is imported or deported, and the DgID. The cluster-shared disk group ID (DgID) is an internal number assigned to the disk group. It can be used in place of the cluster-shared disk group name.

vxdg dginfo

```
vxdg -g<DynamicDiskGroupName> dginfo
```

Returns information about the dynamic disk group that is referenced in `-g<DynamicDiskGroupName>`. Gives the names and numbers of the volumes and the disks in the dynamic disk group. It also includes the dynamic disk group name, its state (either Imported or Deported), its Access level (whether Read/Write or Read-Only) and its dynamic disk group ID. The dynamic disk group ID (DgID) is an internal number assigned to the disk group. It can be used in place of the dynamic disk group name.

Example

```
vxdg -gSalesDynDskGrp dginfo
```

Lists the names and numbers of the volumes and disks in the disk group named "SalesDynDskGrp."

Example of an output for `vxdg dginfo`, when a disk group is set for FastFailover for the disk group QDg1:

On active node

```
C:\vxtool>vxdg -gQDg1 dginfo
```

```
Disk group information for the disk group QDg1
```

```
Name : QDg1
```

```
DiskGroup ID : cc33237e-266b-4b6e-84f6-455191dabdd9
```

```
Disk Group Type : Dynamic (Cluster)
```

```
:
```

```
Status : Imported
```

```
Access : Read/Write
```

```
:
```



```

:
Names of the Volumes are....
\Device\HarddiskDmVolumes\QDG1\QVol1 (Q:)
Names of the Disk are....
Harddisk4
On passive node
C:\>vxdg -g qdgl dginfo
Disk group information for the disk group qdgl
Name : QDG1
DiskGroup ID : cc33237e-266b-4b6e-84f6-455191dabdd9
Disk Group Type : Dynamic (Cluster)
:
Status : Deported
Access : Read-Only
:
:
Names of the Volumes are....
\Device\HarddiskDmVolumes\QDG1\QVol1
Names of the Disk are....
Harddisk4

```

vxdg dginfo (CSDG)

```
vxdg -g<ClusterSharedDiskGroupName> dginfo
```

Returns information about the cluster-shared disk group (CSDG) that is referenced in `-g<ClusterSharedDiskGroupName>`. Gives the names and numbers of the volumes and the disks in the cluster-shared disk group. It also includes the cluster-shared disk group name, its state (either Online or Offline), its Access level (whether Read/Write or not) and its DgID. The cluster-shared disk group ID (DgID) is an internal number assigned to the disk group. It can be used in place of the cluster-shared disk group name.

Example

```
vxdg -gSalesClsShrDskGrp dginfo
```

Lists the names and numbers of the volumes and disks in the cluster-shared disk group named "SalesCIsShrDskGrp."

WRITER NOTES: IS THE FOLLOWING APPLICABLE FOR CSDG?

Example of an output for `vxdg dginfo`, when a cluster-shared disk group is set for FastFailover for the disk group QDg1:

On active node

```
C:\vxtool>vxdg -gQDg1 dginfo

Disk group information for the disk group QDg1

Name : QDg1

DiskGroup ID : cc33237e-266b-4b6e-84f6-455191dabdd9

Disk Group Type : Dynamic (Cluster)

:

Status : Imported

Access : Read/Write

:

:

Names of the Volumes are....

\Device\HarddiskDmVolumes\QDg1\QVol1 (Q:)

Names of the Disk are....

Harddisk4
```

On passive node

```
C:\>vxdg -g qdgl dginfo

Disk group information for the disk group qdgl

Name : QDg1

DiskGroup ID : cc33237e-266b-4b6e-84f6-455191dabdd9

Disk Group Type : Dynamic (Cluster)

:

Status : Deported

Access : Read-Only

:
```

```

:
Names of the Volumes are....

\Device\HarddiskDmVolumes\QDG1\QVol1

Names of the Disk are....

Harddisk4

```

vxdg init

```

vxdg -g<DynamicDiskGroupName> [-f] [-s | -R] [-TLDM] [-o timeout=<n>]
init [<medianame>=<DiskName>... | <p#c#t#l# >...
[basename=<seedname>]

```

Creates a new dynamic disk group with the specified disks.

<code>-g<DynamicDiskGroupName></code>	Identifies the dynamic disk group name. A dynamic disk group name is limited to 18 ASCII characters. It should not have any spaces.
<code>-f</code>	Force the operation.
<code>-s</code>	Creates a cluster dynamic disk group.
<code>-R</code>	Creates a private dynamic disk group, that is a disk group with private dynamic disk group protection.
<code>-TLDM</code>	On Windows Server operating systems: Creates a Microsoft Disk Management Disk Group. Note: Windows Server OS allows only one Microsoft Disk Management Disk Group and automatically assigns the name of the disk group when it is created. Specifying the <code>-TLDM</code> option with the name of the disk group with the <code>-g <diskgroupname></code> attribute results in an error.
<code>-o timeout=<n></code>	Allows the operation to timeout in <code><n></code> seconds.
<code><medianame></code>	Internal disk name that the user specifies. Note: Only ASCII characters are allowed.

<code><DiskName>...</code>	The name of the disk or disks, which the device name (such as <code>Harddisk2</code>) specifies.
<code><p#c#t#l#>...</code>	Another way to indicate a disk name; the #s correspond to port, channel, target, and LUN of the disk.
<code><seedname></code>	A name that when appended with a system-generated number becomes the internal identifier of the disk. Note: Only ASCII characters are allowed.

There can be one or more disks in a dynamic disk group. Disks within a dynamic disk group share a common configuration. Dynamic volumes are created within a dynamic disk group and are restricted to using disks within that group.

You can use this command to create a cluster dynamic disk group, a private dynamic disk group, a Windows Disk Management compatible dynamic disk group, or a Microsoft Disk Management Disk Group.

See [“Disk groups overview”](#) on page 146.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Examples

```
vxdg -gSharedDg1 init -s -TLDM Harddisk0 Harddisk1
```

```
vxdg init -TLDM Harddisk5
```

On Windows Server operating systems, this command creates a Microsoft Disk Management Disk Group. The dynamic disk group includes `Harddisk5`.

```
vxdg -gDG1 init mediaName1=Harddisk1mediaName2=Harddisk2  
Harddisk3 Harddisk4
```

This command creates a new dynamic disk group called "DG1". The dynamic disk group includes `Harddisk1`, `Harddisk2`, `Harddisk3`, and `Harddisk4` and the names `mediaName1` and `mediaName2` are assigned to `Harddisk1` and `Harddisk2` respectively.

vxdg init (CSDG)

```
vxdg -g<ClusterSharedDiskGroupName> [-f] [-s | -r | -csdg] [-TLDM]
[-o timeout=<n>]
```

```
init [<medianame>=<DiskName>... | <p#c#t#l#>...
[basename=<seedname>]
```

Creates a new cluster-shared disk group with the specified disks.

-g<ClusterSharedDiskGroupName>	Identifies the cluster-shared disk group name. A cluster-shared disk group name is limited to 18 ASCII characters. It should not have any spaces.
-f	Force the operation.
-s	Creates a cluster-shared disk group.
-R	Creates a private cluster-shared disk group, that is, a disk group with private cluster-shared disk group protection.
-TLDM	On Windows Server operating systems: Creates a Microsoft Disk Management Disk Group. Note: Windows Server OS allows only one Microsoft Disk Management Disk Group and automatically assigns the name of the disk group when it is created. Specifying the -TLDM option with the name of the disk group with the -g <diskgroupname> attribute results in an error.
-o timeout=<n>	Allows the operation to timeout in <n> seconds.
<medianame>	Internal disk name that the user specifies. Note: Only ASCII characters are allowed.
<DiskName>...	The name of the disk or disks, which the device name (such as Harddisk2) can specify.
<p#c#t#l#>...	Another way to indicate a disk name; the #s correspond to port, channel, target, and LUN of the disk.

<seedname>

A name that when appended with a system-generated number becomes the internal identifier of the disk.

Note: Only ASCII characters are allowed.

There can be one or more disks in a dynamic disk group. Disks within a dynamic disk group share a common configuration. Dynamic volumes are created within a dynamic disk group and are restricted to using disks within that group.

You can use this command to create a cluster dynamic disk group, a private dynamic disk group, a Windows Disk Management compatible dynamic disk group, or a Microsoft Disk Management Disk Group.

See [“Disk groups overview”](#) on page 146.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Examples

```
vxdg -gSharedDg1 init -s -TLDM Harddisk0 Harddisk1
```

```
vxdg init -TLDM Harddisk5
```

On Windows Server operating systems, this command creates a Microsoft Disk Management Disk Group. The dynamic disk group includes Harddisk5.

```
vxdg -gDG1 init mediaName1=Harddisk1mediaName2=Harddisk2  
Harddisk3 Harddisk4
```

This command creates a new dynamic disk group called "DG1". The dynamic disk group includes Harddisk1, Harddisk2, Harddisk3, And Harddisk4 and the names mediaName1 and mediaName2 area that is assigned to Harddisk1 and Harddisk2 respectively.

vxdg adddisk

```
vxdg -g<DynamicDiskGroupName> [-o timeout=<n>] adddisk  
[<medianame>=<DiskName>... | <p#c#t#l#>... [basename=<seedname>]
```

Adds one or more basic disks that are referenced in <DiskName>... | <p#c#t#l#>... to the dynamic disk group that is referenced in -g<DynamicDiskGroupName>.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	The name of the disk group that the disk is added to.
<code>-o timeout=<n></code>	Allow the operation to timeout in <code><n></code> seconds.
<code><medianame></code>	Internal disk name that the user specifies. Note: Only ASCII characters are allowed.
<code><DiskName>...</code>	The name of the disk or disks, which the device name (such as <code>Harddisk2</code>) specifies.
<code><p#c#t#l#>...</code>	Another way to indicate a disk name; the <code>#</code> s correspond to port, channel, target, and LUN of the disk.
<code><seedname></code>	A name that when appended with a system-generated number becomes the internal identifier of the disk. Note: Only ASCII characters are allowed.

The ellipses indicate that you can specify more than one disk.

You can no longer create a dynamic disk group implicitly by using a different dynamic disk group name. Dynamic disk group creation and all the options that are associated with the operation are now done through `vxdg init` instead.

Examples

```
vxdg -gSalesDynDskGrp adddisk Harddisk4
```

Adds `Harddisk4` to `SalesDynDskGrp`.

```
vxdg -gclustr_dyndskgrp1 adddisk Harddisk5
```

Adds `Harddisk5` to the cluster dynamic disk group `clustr_dyndskgrp1`.

vxdg rmdisk

```
vxdg [-g<DynamicDiskGroupName>] [-f] rmdisk <DiskName> | p#c#t#l#>
```

Reverts a disk that is referenced in `<DiskName> | p#c#t#l#>` from a dynamic disk to a basic disk. `<DiskName>` can be the disk name or path name of the device, such as `Harddisk2` or `\Device\Harddisk2`. The disk must not contain any volumes when it is reverted.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	The name of disk group that the disk is removed from.
<code>-f</code>	Forces the command; use this option with caution.
<code><DiskName></code>	The disk name of the disk being reverted, which the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>) specifies. The internal disk name must be preceded by the <code>-g</code> option (for example, <code>-gDG2 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk name; the <code>#</code> s correspond to port, channel, target, and LUN of the disk.

Example

```
vxdg rmdisk Harddisk2
```

Reverts `Harddisk2` to a basic disk.

vxdg import

```
vxdg -g<DynamicDiskGroupName> [-n<NewDynamicDiskGroupName>] [-s |-d] [-f] [-C] [-o] timeout=<n> import
```

Imports the dynamic disk group that is referenced in `-g<DynamicDiskGroupName>` for use on this computer. You can use its name from the original computer or rename it with `-n<NewDynamicDiskGroupName>`.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg import` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	The name of the disk group that is imported.
---	--

<code>-s</code>	Imports the disk group as a cluster disk group. Using this command with the <code>-s</code> option is the only way you can convert a regular secondary disk group to a cluster disk group. With SFW, it is only required to use the <code>-s</code> the first time you create the cluster disk group.
<code>-d</code>	Converts a cluster disk group to a regular disk group.
<code>-f</code>	Forces the command; use this option with caution. The <code>-f</code> option can be used to force the import of a cluster disk group when a minority of disks in the disk group is available. See “Unable to bring a cluster disk group online when a minority of the disks in the disk group are available” on page 555.
<code>-C</code>	Clears the disk group’s host ID. See more about the host ID below.
<code>-o timeout=<n></code>	The <code>-o timeout=<n></code> option allows the operation to timeout in <code><n></code> seconds. This is useful for scripting purposes.

About a disk group host ID:

When a dynamic disk group is created or imported, all disks in the dynamic disk group are stamped with the computer’s host ID. Typically, a dynamic disk group cannot be imported if any of its disks are stamped with a non-matching host ID. This mechanism provides a check in cases when disks can be accessed from more than one host.

In Storage Foundation for Windows, if a dynamic disk group is in use by one computer when another computer attempts to import this dynamic disk group again, the import fails because the host IDs of these two computers are different.

If a dynamic disk group is successfully deported, the host id is cleared and the group can be safely imported to another computer. With SFW, you do not need the `-s` option. A cluster disk group remains a cluster disk group when it is imported. If you want to change a cluster disk group back to an ordinary disk group, you deport it and then import it with the `-d` option.

To move a dynamic disk group from one computer to another when the dynamic disk group was not cleanly deported (for example, if the host computer crashed),

you must specify the `-c` option to clear the original host id and stamp a new host id onto the dynamic disk group.

Note: While importing a cluster disk group as a cluster-shared disk group (CSDG), you must use the `-c` option.

Example

```
vxdg -gDynDskGrp1 -nAcctgDynDskGrp import
```

Imports the dynamic disk group previously known as "DynDskGrp1" and renames it "AcctgDynDskGrp."

vxdg import (CSDG)

```
vxdg -g<DiskGroupName> [-n<NewDiskGroupName>] [-s | -d | -csdg] [-c] [-f] [-o timeout=<n>] import
```

Imports a dynamic disk group or cluster disk group that is referenced in `-g<DiskGroupName>` as a cluster-shared disk group (CSDG). You can use its original name or rename it using the `-n<NewDiskGroupName>` option.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg import` command.

The following attributes apply:

<code>-g<DiskGroupName></code>	The name of the disk group that you want to import.
<code>-s</code>	Imports the disk group as a cluster disk group. Using the command with the <code>-s</code> option is the only way you can convert a regular disk group to a cluster disk group. With SFW, it is only required to use the <code>-s</code> the first time you create the cluster disk group.
<code>-d</code>	Converts a cluster disk group to a regular disk group.
<code>-csdg</code>	Converts a dynamic or cluster disk group to a cluster-shared disk group (CSDG).

<code>-c</code>	Clears the disk group's host ID. See more information about the host ID below.
<code>-f</code>	Forces the command; use this option with caution. The <code>-f</code> option can be used to force the import of a disk group when only a minority of disks in the disk group is available. See "Unable to bring a cluster disk group online when a minority of the disks in the disk group are available" on page 555.
<code>-o timeout=<n></code>	The <code>-o timeout=<n></code> option allows the operation to timeout in <code><n></code> seconds. This is useful for scripting purposes.

Example:

```
vxdg -gDynDskGrp1 -nClsShrDskGrp1 import
```

Imports the dynamic disk group previously known as "DynDskGrp1" as a cluster-shared disk group and renames it "ClsShrDskGrp1".

About the disk group host ID:

When a dynamic disk group is created or imported, all the disks in the dynamic disk group are stamped with the computer's host ID. Typically, a dynamic disk group cannot be imported if any of its disks are stamped with a non-matching host ID. This mechanism provides a check in cases when disks can be accessed from more than one host.

In Storage Foundation for Windows, if a dynamic disk group is in use by one computer when another computer attempts to import this dynamic disk group again, the import fails because the host IDs of these two computers are different. If a dynamic disk group is successfully deported, the host ID is cleared and the group can be safely imported to another computer. With SFW, you do not need the `-s` option.

A cluster disk group remains a cluster disk group when it is imported. If you want to change a cluster disk group back to an ordinary disk group, you deport it and then import it with the `-d` option.

To move a dynamic disk group from one computer to another when the dynamic disk group was not cleanly deported (for example, if the host computer crashed), you must specify the `-c` option to clear the original host ID and stamp a new host ID onto the dynamic disk group.

Note: While importing a cluster disk group as a cluster-shared disk group (CSDG), you must use the `-c` option.

vxdg deport

```
vxdg -g<DynamicDiskGroupName> [-f] deport
```

Deposits the dynamic disk group that is referenced in `-g<DynamicDiskGroupName>` in preparation for its being imported to another computer. Disks and volumes cannot be accessed until the dynamic disk group is imported. The `-f` option forces the deport of the disk group if one or more of its volumes are still in use. Storage Foundation for Windows does not let you deport a cluster disk group that is a VCS or Microsoft Failover Clustering disk group resource.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg deport` command.

Example

```
vxdg -gDynDskGrp1 deport
```

Deposits the dynamic disk group named "DynDskGrp1."

vxdg destroy

```
vxdg -g<DynamicDiskGroupName> [-f] destroy
```

Deletes the dynamic disk group that is referenced in `-g<DynamicDiskGroupName>`. The `-f` option forces the delete of the disk group.

Example

```
vxdg -gDynDskGrp1 destroy
```

Deletes the dynamic disk group named "DynDskGrp1."

vxdg protect

```
vxdg -g<DynamicDiskGroupName> protect
```

Adds private dynamic disk group protection to a secondary dynamic disk group on the current host. Private dynamic disk group protection allows administrators to partition storage and assign exclusive ownership in a SAN or shared disk array.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg protect` command.

Private dynamic disk group protection uses hardware locking techniques to protect a secondary dynamic disk group that is located on shared storage from access by other hosts that are connected to the shared storage pool. The hardware locking is implemented by using a SCSI reservation thread to maintain a current reservation for each disk in a protected group.

If disks are added to a dynamic disk group that has a SCSI reservation, they get a reservation automatically.

As long as the dynamic disk group remains on the same host, the protection is enabled, even through restarts. If the dynamic disk group is deported and imported on another host, the private dynamic disk group protection is removed and the command must be redone to restore the protection.

Also, if you uninstall the storage management software, the dynamic disk group protection information is removed. If you then reinstall the storage management software and want to restore the dynamic disk group protection, you must redo the command.

See [“Partitioned shared storage with private dynamic disk group protection”](#) on page 162.

Example

```
vxdg -gDynDskGrp2 protect
```

Adds private dynamic disk group protection to DynDskGrp2.

vxdg release

```
vxdg -g<DynamicDiskGroupName> release
```

Removes private dynamic disk group protection.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg release` command.

Example

```
vxdg -gDynDskGrp2 release
```

Removes private dynamic disk group protection from DynDskGrp2.

vxdg upgrade

```
vxdg -g<DynamicDiskGroupName> [-T <version>] upgrade
```

Upgrades a dynamic disk group that was originally created on an earlier version of Volume Manager or its related programs to the current Storage Foundation for Windows dynamic disk group version. The optional `-T <version>` parameter lets you upgrade a disk group that was created under Disk Management or earlier versions of Volume Manager for Windows to the Volume Manager disk group version.

In most cases, you'll want to upgrade an existing disk group so that it can take advantage of the dynamic disk group capabilities of Storage Foundation for Windows.

If you don't plan to make changes to these existing disk groups, you do not have to upgrade them; upgrading them lets you use additional features that affect a disk group, such as SFW's support for VSS, the Volume Shadow Copy Service feature. SFW incorporates VSS support into FlashSnap. If you plan to use the FlashSnap option with these existing disk groups, then you should upgrade them to the current disk group type.

Note: If you upgrade a disk group to SFW, you cannot import it on another server that is running earlier versions of Volume Manager or Disk Management. Once a disk group version is upgraded, it cannot be changed back to an earlier disk group version.

For information about the supported disk group versions, See "[Dynamic disk group properties](#)" on page 165.

Note: SFW does not let you upgrade a dynamic disk group that contains a boot or system volume. It must remain at the Disk Management disk group version level to allow for the fact that the program could be uninstalled and that a system and boot disk group would become inaccessible if it were a higher disk group version.

Examples

```
vxdg -gDynDskGrp1 upgrade
```

This command upgrades Storage Foundation for Windows dynamic disk group `DynDskGrp1` to the latest version.

```
vxdg -gDynDskGrp2 -T60 upgrade
```

This command upgrades Storage Foundation for Windows dynamic disk group `DynDskGrp2` to the version 60.

```
vxdg -gDynDskGrp2 -TSFW upgrade
```

This command can be used to convert Microsoft Disk Management Disk Group into SFW dynamic disk group. Windows Disk Management does not manage a converted disk group. It is only accessible and managed by SFW. The conversion cannot be undone.

Note: Converting a Microsoft Disk Management Disk Group that was created using iSCSI disks to an SFW dynamic disk group is currently not supported.

You need to complete the conversion procedure with the following:

- Restart the server.
- Import the converted disk group using the `vxdg import` command. Specify the `-c` option in the command to clear the host ID.
- If advanced disk group features are needed, upgrade the disk group version using the `vxdg upgrade` command.

vxdg repldisk

```
vxdg [-g<DynamicDiskGroupName>] repldisk <"Missing Disk (disk#)" | DiskName |p#c#t#l# >=<SpareDiskName>
```

Replaces a missing disk that is specified by `"Missing Disk (disk#)" | DiskName |p#c#t#l# >` moving all the subdisks to the disk that is specified by `<SpareDiskName>`. The disk specified by `<SpareDiskName>` must be a basic blank disk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdg repldisk` command.

The following attributes apply:

`-g<DynamicDiskGroupName>`

This parameter is required only if the internal disk name is used to identify the missing disk.

`"Missing Disk (disk#)"`

Specifies the missing disk. This string represents how the program identifies the missing disk. Because the string `"Missing Disk (disk#)"` has spaces, you must use quotation marks around it. The program specifies the `(disk#)`, such as `(disk2)`.

<i>DiskName</i>	The internal disk name of the missing disk. Requires the use the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk name; the <code>#s</code> correspond to port, channel, target, and LUN.
<code><SpareDiskName></code>	The name of the replacement disk.

This command is listed under dynamic disk group commands because it applies only to dynamic disks.

Example

```
vxdg repldisk "Missing Disk (disk2)"=Harddisk3
```

This command moves the contents of a missing dynamic disk that is designated as "Missing Disk (disk2)" to a basic disk, which in this case is Harddisk3. In the process, Harddisk3 becomes a dynamic disk. Active redundant volumes (mirrored and RAID-5) on the replaced disk is resynchronized automatically after disk replacement.

```
vxdg -gDG1 repldisk Disk2=Harddisk3
```

This is an example of the same command using the internal disk name parameter.

vxdg split

```
vxdg -g<SourceDynamicDiskGroupName> -n<TargetDynamicDiskGroupName>
[-i] [-f] [-s|-d] [-y] [-v] split [<Object>...]
```

Splits a dynamic disk group into two dynamic disk groups. The disks and volumes that are in the split-off disk group must be healthy. Also, the volumes that are in the new disk group must be complete for the command to succeed.

Microsoft Disk Management Disk Groups do not support the `vxdg split` command.

See [“Dynamic disk group split”](#) on page 274.

<code>-g<SourceDynamicDiskGroupName></code>	Name of the current dynamic disk group.
<code>-n<TargetDynamicDiskGroupName></code>	Name of the new dynamic disk group.
<code>-i</code>	Makes the new dynamic disk group to be in the Imported state.

-f	Forces the split. Use this command with care.
-s	Makes the new dynamic disk group a cluster dynamic disk group.
-d	Changes the target disk group from a cluster disk group to a regular disk group.
-y	Query for the split closure. See the first example below.
-v	Splits all disks in the dynamic disk group that contain snapshot volumes. Do not use the <i><Object></i> parameter when you use this option.
<i><Object></i>	Name of a disk or volume, or a drive letter. You can split by either disks or volumes. This option is not needed when using the <i>-v</i> option. The internal volume and disk name can be used. The <i><DrivePath></i> parameter is supported for a volume name. Also, a disk can be indicated by the <i>p#c#t#l#</i> parameter; the <i>#s</i> correspond to port, channel, target, and LUN.

A primary dynamic disk group that is created in Storage Foundation for Windows cannot be split because it contains a boot or system drive. If you have dynamic disk groups created in an early release of Volume Manager, you need to upgrade them to be compatible with Storage Foundation for Windows dynamic disk groups before you can use the `vxdg split` and `join` commands. To do this you use the `vxdg upgrade` command.

By default, the new dynamic disk group is in the Deported state after the split. You can use the `-i` option to have it be in the Imported state. The source dynamic disk group remains in the Imported state.

With the VEA GUI split command, if the source disk group is a cluster disk group or a disk group with private dynamic disk group protection, the resulting target disk

group is of the same disk group type. With a CLI split command for either of these two disk group types, you must use the `-i` switch. In a CLI split operation with either of these disk group types, the target disk group does not assume the disk group type of the source disk group until after the target disk group is imported.

See [“Dynamic Disk Group Split and Join”](#) on page 272.

Note: In earlier versions of Volume Manager, it was necessary to designate the `-s` switch with every command involving a cluster disk group. However, in SFW, you use the `-s` switch only the first time you designate a cluster disk group. The `-d` switch has been added to allow changing a cluster disk group to an ordinary disk group.

Before using dynamic disk group split and join commands with a cluster on shared storage, you should review the considerations for doing so.

See [“Using Dynamic Disk Group Split and Join with a cluster on shared storage”](#) on page 282.

Examples

```
vxdg -gDynDskGrp1 -y -nDynDskGrp2 split Harddisk5 Harddisk7
```

This command makes a query to determine whether Harddisk5 and Harddisk7 in a dynamic disk group named "DynDskGrp1" comprise the total disks that are needed to have a dynamic disk group split where all the volumes in the split-off dynamic disk group are complete. Here is the output from the command:

```
Following is the Closure for the requested Splitoperation.
Names of the Disks are....
Harddisk7
Harddisk5
Harddisk6
Names of the Volumes are....
\Device\HarddiskDmVolumes\DynDskGrp1\Volume1
\Device\HarddiskDmVolumes\DynDskGrp1\Stripe1
```

The output indicates that to have a successful split, or what is called "split closure," Harddisk6 must be added. The example of the command to perform the actual split is as follows:

```
vxdg -gDynDskGrp1 -i -nDynDskGrp2 split Harddisk5 Harddisk6 Harddisk7
```

This command results in successfully splitting the dynamic disk group DynDskGrp1 with the target dynamic disk group DynDskGrp2 in the Imported state. The new dynamic disk group has the disks Harddisk5, Harddisk6, and Harddisk7.

The example that follows designates the volumes to be included in a new target disk group. Note the path name that is needed for volumes.

```
vxdg -gDynDskGrp1 -i -nDynDskGrp2 split  
\Device\HarddiskDmVolumes\DynDskGrp1\mirrorvol1  
\Device\HarddiskDmVolumes\DynDskGrp1\mirrorvol2
```

This command results in successfully splitting the dynamic disk group DynDskGrp1 with the target dynamic disk group DynDskGrp2 in the Imported state. The new dynamic disk group contains the volumes \Device\HarddiskDmVolumes

\DynDskGrp2\mirrorvol1 and \Device\HarddiskDmVolumes

\DynDskGrp2\mirrorvol2.

```
vxdg -gtest -nnew split plc0t1l0 plc0t2l0
```

This command performs the split operation. The name of the new disk group is "new." This command shows an example of the use of the `p#c#t#l#` with the `<Object>` parameter.

```
vxdg -gtest -nnew split Disk1 Disk2
```

This command performs the split operation. This example shows the use of the internal disk name with the `<Object>` parameter.

```
vxdg -gtest -nnew split Volume1 Volume2
```

This command performs the split operation. This example shows the use of the internal volume name with the `<Object>` parameter.

```
vxdg -gtest -nnew split L:\Folder
```

This command performs the split operation. This example shows the use of `<DrivePath>` with the `<Object>` parameter.

```
vxdg -gtest -nnew split Volume1 Volume2 -f
```

This command performs the split operation. This example shows the use of the `-f` force option.

vxdg recover

```
vxdg -g<DynamicDiskGroupName> [-s|-d] recover
```

Recovers a dynamic disk group that fails because of a system crash or other problem during a dynamic disk group split operation. This command is used for manual recovery of a disk group when the disk group requires such recovery. The `-s` option

recovers the disk group as a cluster disk group in the Imported state. The `-d` parameter is used to change a cluster disk group to a regular disk group.

More information is available about automatic and manual modes of recovery.

See [“Recovery for the split command”](#) on page 277.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support cluster disk groups.

Example

```
vxdg -gDynDskGrp2 recover
```

Recovers DynDskGrp2 after it failed because of a system crash during a dynamic disk group split operation.

vxdg join

```
vxdg -g<SourceDynamicDiskGroupName> -n<TargetDynamicDiskGroupName>  
[-C] [-P] [-o timeout=<n>] join
```

Joins two dynamic disk groups together. If some disks are missing or are not functioning, the join command fails.

Note: Microsoft Disk Management Disk Groups do not support the `vxdg join` command.

`-g<SourceDynamicDiskGroupName>`

Name of the source dynamic disk group, the dynamic disk group that is added to the target dynamic disk group. The source dynamic disk group ceases to exist after the join.

`-n<TargetDynamicDiskGroupName>`

Name of the target dynamic disk group, the dynamic disk group that exists after the join.

-C	Clear host ID. This parameter should be used when you want to import a disk group from another node and that disk group was not properly deported. A host id might not be cleared, for example, if the host computer became nonfunctional. When a host id has not been cleared, SFW does not permit the join without the Clear host ID override. Use this override with care.
-P	Allow Partial Join. See the explanation of the PartialJoin command: See “About partial disk group join” on page 281.
-o timeout=<n>	Allow the operation to timeout in <n> seconds. This parameter is useful for scripting.

The disk group type after the join is the type of the target disk group. For example, if the target disk group before the join had private dynamic disk group protection, the resulting disk group has private dynamic disk group protection after the join. Also, if a cluster software, such as VCS or Microsoft Failover Clustering, monitors the disk group, a cluster dynamic disk group is not allowed to be the source group for a join command. However, a cluster disk group can be a target dynamic disk group in a join command.

See [“Dynamic Disk Group Split and Join”](#) on page 272.

Example

```
vxdg -gDynDskGrp2 -nDynDskGrp1 join
```

This command joins DynDskGrp2 back into DynDskGrp1. You can also join dynamic disk groups together that were not originally split apart.

```
vxdg join -gddg -ncdg -o timeout=200 join
```

This command joins the dynamic disk group ddg to dynamic disk group cdg. The example shows the use of the `-o timeout=<n>` parameter.

vxdg reclaim

```
vxdg [-g<DynamicDiskGroupName>] [-b] reclaim  
option=<volumes|freespaces|all>
```

Reclaim storage space from thin provisioned disks in a dynamic disk group.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group.
<code>-b</code>	Run the command in the background.
<code>volumes</code>	Reclaim storage space from volumes.
<code>freespaces</code>	Reclaim free space from disks.
<code>all</code>	Reclaim storage space from volumes and free space from disks. (Default option.)

Note: The operation to reclaim storage from thin provisioned disks is not supported for RAID-5, DCO, or Volume Replicator DCM volumes.

Note: The reclaim storage operation is disabled during a shrink volume operation. The reclaim storage operation is enabled when the shrink volume operation has completed.

Example

```
vxdg -gDG1 reclaim option=freespaces
```

This command reclaims only the freespace of the thin provisioned disks in the DG1 dynamic disk group.

vxdg reclaimcancel

```
vxdg [-g<DynamicDiskGroupName>] reclaimcancel  
option=<volumes|freespaces|all>
```

Immediately cancels the reclaim operation. Only a portion of the storage space of the dynamic disk group may have been reclaimed.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group.
<code>volumes</code>	Reclaim storage space from volumes.
<code>freespaces</code>	Reclaim freespace from disks.
<code>all</code>	Reclaim storage space from volumes and freespace from disks. (Default option.)

Example

```
vxdg -gDG1 reclaimcancel
```

This command immediately cancels the reclaim operation for volumes and freespace of the thin provisioned disks in the DG1 dynamic disk group. (No option is specified, so the default option of volumes and freespace is used.)

vxdg latestart

```
vxdg -g<DynamicDiskGroupName> latestart on|off
```

The following attributes are:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group that is made available after system startup.
<code>on</code>	Enables the specified dynamic disk group to be imported after system startup.
<code>off</code>	Disables the specified dynamic disk group from being imported after system startup.

Generally, dynamic (non-cluster) disk groups are automatically imported during system startup. However, some types of storage are not available during the system startup which would cause the import to fail. In a clustered environment, disk groups are imported by the cluster application and do not need to have `vxdg latestart` enabled.

Specifying `on` in the `vxdg latestart` command enables the dynamic disk group that is referenced in `-g<DynamicDiskGroupName>` to be imported after system startup by the Veritas DG Delayed Import Service (VxDgDI). VxDgDI can import the dynamic disk group after it is made dependent on the service that controls the

storage. This allows the required time for the storage to become available. Applications that rely on storage that is imported by the VxDgDI service may also need to be made dependent on VxDgDI so that they may proceed when their storage is available. Making the VxDGDI service startup type automatic allows the service to run every time the system is started.

Note: To configure a dependency for iSCSI, use VxDgDI to import the dynamic disk group after it is made dependent on the iSCSI service (MSiSCSI) that controls the storage.

As long as the dynamic disk group remains on the same host, `vxdg latestart` is enabled, even through restarts. If the dynamic disk group is deported and imported on another host, `vxdg latestart` must be reenabled on the new host.

Note: For SFW, newly created dynamic disk groups that contain iSCSI disks, or existing dynamic disk groups that have iSCSI disks added to them, are automatically configured to be imported after system startup. Using the `vxdg latestart` command is not required for these dynamic disk groups.

To set up `vxdg latestart`:

- 1 Edit the Windows registry to make VxDgDI dependent on the service that controls the storage.
- 2 Enter the `vxdg latestart` command at the command line.

For example

```
vxdg -gDynDskGrp2 latestart on
```

enables the dynamic disk group "DynDskGrp2" to be imported after system startup.

- 3 Applications that rely on storage that is imported by the VxDgDI service have their storage available automatically after the VxDgDI service completes its startup process. However applications that start as a Windows service need to be made dependent on the VxDgDI service with the Windows Registry Editor before their storage is available.

vxdg refreshff

```
vxdg -g<diskgroupname> refreshff
```

This command refreshes the state of the disk groups on a system to support the new Deported Read-Only disk group state of fast failover.

Sometimes, the disk groups on passive nodes do not reflect the new Read-Only state, even though FastFailover is enabled for the VMDg resources. This typically occurs in cases where a new node is added to a service group where FastFailover is already enabled. This command must be run on the newly added node for it to support the fast failover disk group state.

vxclus

The `vxclus` utility makes it possible to bring a Microsoft Failover Clustering disk group online on a node with a minority of the disks in the disk group. The `vxclus` utility creates an entry in the Registry that enables the cluster resource to be brought online. Once `vxclus enable` is executed, you can bring the resource online with the Windows Server's Failover Cluster Manager snap-in.

With SFW, the `vxclus` utility has added keywords that allow a cluster disk group to reside on the same bus as the cluster node's system and boot disk. These keywords are necessary to support booting from a SAN. Also your hardware array must also have support for booting from a SAN. Refer to your hardware array manufacturer's instructions and the Microsoft Web site for more information about booting from a SAN.

Note: The `vxclus` commands are needed only with Microsoft Failover Clustering.

Bringing cluster disk groups online on a node with a minority of the disk group's disks is described in more detail in the following topic:

See [“Unable to bring a cluster disk group online when a minority of the disks in the disk group are available”](#) on page 555.

<code>enable</code>	Allows a cluster disk group to be brought online when a minority of disks in the disk group is available.
<code>disable</code>	Disables the capability to bring a cluster disk group online when a minority of disks in the disk group are available.
<code>cleanup</code>	Disables for all cluster disk groups on a node the capability to bring the disk group online when a minority of disks in the disk group is available. The cleanup option is not applicable to the <code>UseSystemBus ON OFF</code> support.
<code>UseSystemBus ON</code>	Sets up the conditions to allow a cluster disk group to be created that can reside on the same bus as the cluster node's system or boot disk.

UseSystemBus OFF Removes the capability of a creating a cluster disk group that can reside on the same bus as the cluster node's system or boot disk.

Typing the following sequence for each keyword brings up a description of its syntax:

```
vxclus <Keyword> -?
```

Warning: To avoid possible data corruption, make sure that a majority of a disk group's disks do not become available and online on one node while another node with a minority of the disk group's disks also has the disk group online.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxclus` command.

vxclus enable

```
vxclus enable -g<DynamicDiskGroupName> [-p]
```

Enables a designated cluster disk group for force import so that it may be brought online when a minority of disks in the disk group are available. The `vxclus` utility creates an entry in the Windows registry that enables the cluster resource for force import. Once `vxclus enable` is executed, you can bring the resource online with Windows Server's Failover Cluster Manager snap-in. After the resource is brought online, the `vxclus` force import functionality is disabled. However if `-p` is specified, the entry that is made in the Windows registry is such that the `vxclus` force import functionality remains enabled. This allows persistent force import of the designated cluster disk group so that this resource can always be brought online with the Failover Cluster Manager snap-in.

The following attributes apply:

<code>-g<SourceDynamicDiskGroupName></code>	Name of the cluster disk group to be enabled for force import.
<code>-P</code>	Enables persistent force import of the cluster disk group.

Example

```
vxclus enable -gDG2 -p
```

Enables the DG2 cluster disk group to be brought online on a node when a minority of disks in the disk group are available. Also the DG2 clusterdisk group is enabled

so that it is always ready to be brought online with Windows Server's Failover Cluster Manager snap-in.

vxclus disable

```
vxclus disable -g<DynamicDiskGroupName>
```

Disables the capability to bring the designated cluster disk group online on a node when a minority of disks in the disk group are available. This command is necessary only if you have executed the `vxclus enable` command but have not yet brought the disk group online with Windows Server's Failover Cluster Manager snap-in and later decide that you do not want to bring the disk group online if there are not many disks available in a disk group.

Example

```
vxclus disable -gDG2
```

Removes from the DG2 cluster disk group the capability to be brought online when a minority of disks in the disk group are available.

vxclus cleanup

```
vxclus cleanup
```

Disables for all cluster disk groups on a node the capability to bring a disk group online when a minority of disks in the disk group are available. This command is necessary only if you have executed the `vxclus enable` command on multiple cluster disk groups but you have not yet brought the cluster disk groups online with Windows Server's Failover Cluster Manager snap-in and you decide that you do not want to bring these disk groups online.

Example

```
vxclus cleanup
```

On the current cluster node, this command turns off the `vxclus` capability for any cluster disk groups that had this capability enabled.

vxclus UseSystemBus ON

```
vxclus UseSystemBus ON
```

This command makes the registry changes necessary to allow a dynamic cluster disk group to be created on the same bus as the cluster node's system or boot disk. In previous versions of Volume Manager, the program did not allow this to happen. The reason for making this command available is for users who want to have the

boot or system disk and the cluster disk group located on a SAN to boot from the SAN.

Warning: If you decide that you want to have a cluster disk group on the same bus as the computer node's system or boot disk and you want to boot from a SAN, then you need to follow the specific instructions of your array manufacturer for booting from a SAN. See the Microsoft Windows Server Knowledge Base article 305547 for details of the considerations involved when booting from a SAN.

Example

```
vxclus UseSystemBus ON
```

Makes it possible to create a cluster disk group on the same bus as the cluster node's system or boot disk.

vxclus UseSystemBus OFF

```
vxclus UseSystemBus OFF
```

Reverses the `vxclus UseSystemBus ON` command, removing the Registry changes that make it possible to create a cluster disk group that can reside on the same bus as the cluster node's system or boot disk.

Example

```
vxclus UseSystemBus OFF
```

Removes the capability of creating a cluster disk group on the same bus as the cluster node's system or boot disk.

vxdisk

The `vxdisk` utility has commands relating to disks.

<code>diskinfo</code>	Gives information about a disk.
<code>attach</code>	Connects the disk and enables access to the disk.
<code>detach</code>	Stops disk access and disconnects the disk.
<code>list</code>	Gives information about all the disks on a managed server or on a specified disk group.
<code>list (CSDG)</code>	Gives information about all the disks on a managed server or on a specified cluster-shared disk group (CSDG).

<code>merge</code>	Merges a foreign disk.
<code>reactivate</code>	Reactivates a disk (restarts a disk).
<code>sig</code>	Places a disk signature on a disk.
<code>convertdisk</code>	Converts the partitioning scheme of a disk to MBR or GPT.
<code>set track</code>	Sets track alignment offset for data storage.
<code>trackaligninfo</code>	Displays track alignment setting for host and track alignment offset and settings for all arrays.
<code>set hotreloc</code>	This option sets hot relocation at the host or disk level.
<code>setsite</code>	This option is used to add disks to a particular site.
<code>rmsite</code>	This option is used to remove disk or disk from a site.
<code>ping</code>	Identifies a physical disk.
<code>offline</code>	
<code>online</code>	
<code>reclaim</code>	Reclaim storage space from a thin provisioned disk.
<code>reclaimcancel</code>	Immediately cancel reclaim operation

Each keyword is followed by `<DiskName>`, the name of the disk, which is indicated by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The internal disk name must be preceded by the `-g` disk group option (for example, `-gDG1 Disk2`). Another way to specify a disk is with the `p#c#t#l#` parameter. The `#s` correspond to port, channel, target, and LUN of the disk.

Typing the following sequence for each keyword brings up a description of its syntax:

```
vxdisk <Keyword> -?
```

vxdisk diskinfo

Provides information on the disk that is referenced in `<DiskName>`. The information includes the disk name, the dynamic disk group with which the disk is associated, capacity (indicated as *Length*), disk partition style, free space, device type, port, channel, target ID, LUN, the disk signature, serial number, track alignment settings, Thin Provisioning type (ThinNoReclaim or ThinReclaim for thin provisioned disks), etc. If the disk has associated subdisks, they are listed.

```
vxdisk [-g <DynamicDiskGroupName>] diskinfo [-e] <DiskName>
```

The following attributes apply:

<code>-e</code>	Is an optional parameter. Use the <code>-e</code> option if you want to display extended attribute information for arrays in addition to the basic disk information. Extended attribute information available for display depends on the array type. It can include vendor ID, product ID, revision ID, cabinet serial number, array volume ID, array LUN type, array RAID level, array snapshot LUN, array replication LUN, array media type, array transport protocol, array port WWN number, array port serial number, array controller ID, and array hardware mirror.
<code>DiskName</code>	Name of the disk, e.g., <code>Harddisk2</code> ; or internal disk name (requires use of <code>-g</code> option), e.g., <code>-g MyDg Disk2</code> The <code>-g</code> option is needed only if you use the internal disk name to indicate the disk name (for example, <code>-gDG1 Disk2</code>). If the disk is site tagged, then it displays the site information if the disk name is specified. See “Administering disks for site-based allocation” on page 454.
<code>p#c#t#l#</code>	Corresponds to the Port, Channel, Target, LUN of a disk.

Note: Not all manufacturers of thin provisioned disks enable the display of the current provisioned size.

Example

```
vxdisk diskinfo Harddisk2
```

Gives information about `Harddisk2`.

vxdisk attach

```
vxdisk [-g<DynamicDiskGroupName>] attach <DiskName | p#c#t#l# >
```

Attaches a disk that was detached.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk attach` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	This parameter is required only if the internal disk name is used.
<code><DiskName></code>	The internal disk name of the missing disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of a disk.

Example

```
vxdisk -gDG1 attach Disk2
```

Attaches a previously detached disk, Disk2, to disk group DG1.

vxdisk detach

```
vxdisk [-g<DynamicDiskGroupName>] detach <DiskName | p#c#t#l# >
```

Detaches a disk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk detach` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	This parameter is required only if the internal disk name is used.
<code><DiskName></code>	The internal disk name of the missing disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of a disk.

Example

```
vxdisk -gDG1 detach Disk2
```

Detaches a previously attached disk, -gDG1 Disk2.

vxdisk list

Lists the device name, internal disk name, disk group, disk partition style, size, free space, LUN, status, serial number, ThinProvisioningType (ThinNoReclaim or ThinReclaim for thin provisioned disks), ProvisionedSize (current size of thin provisioned disks), Site, etc. for all the disks on the managed server.

See [“Administering disks for site-based allocation”](#) on page 454.

```
vxdisk [-g <DynamicDiskGroupName>] [-v] list
```

The following attributes apply:

- g The -g option limits the list of disks to the specified disk group.
- v The -v (verbose) option displays signature, status, bus type, port, target, channel, LUN, track alignment offset, and track alignment setting information for the disks in the list.

Note: Not all manufacturers of thin provisioned disks enable the display of the current provisioned size.

Example

```
vxdisk -gDB1 -v list
```

Gives the information about all the disks in the DB1 disk group. The -v option includes the signature, status, bus type, port, channel, target, and LUN information.

vxdisk list (CSDG)

Lists the device name, internal disk name, cluster-shared disk group, disk partition style, size, free space, LUN, status, serial number, ThinProvisioningType (ThinNoReclaim or ThinReclaim for thin provisioned disks), ProvisionedSize (current size of thin provisioned disks), Site, etc. for all the disks on the managed server.

See [“Administering disks for site-based allocation”](#) on page 454.

```
vxdisk [-g <ClusterSharedDiskGroupName>] [-v] list
```

The following attributes apply:

- g The -g option limits the list of disks to the specified cluster-shared disk group.
- v The -v (verbose) option displays signature, status, bus type, port, target, channel, LUN, track alignment offset, and track alignment setting information for the disks in the list.

Note: Not all manufacturers of thin provisioned disks enable the display of the current provisioned size.

Example

```
vxdisk -gDB1 -v list
```

Gives the information about all the disks in the DB1 cluster-shared disk group. The -v option includes the signature, status, bus type, port, channel, target, and LUN information.

vxdisk merge

```
vxdisk -g<DynamicDiskGroupName> merge <DiskName | p#c#t#l#>
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk merge` command.

In the GUI, this command is called Merge Foreign Disk. If you remove a disk from the server and also remove it in Storage Foundation for Windows and then decide to reattach it to the server as a member of the same dynamic disk group, you need to use this command. The command reinstates the disk to its old status as a member of its former dynamic disk group on that server.

You also need to use this command if you remove a disk without taking its disk group offline, move the disk to another server, and then attempt to reinstall it in the original server. The command is necessary in this case because the disk has the other server's disk group ID.

See [“Merge foreign disk command”](#) on page 538.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group.
<code><DiskName></code>	The designated hard disk, which can be specified by the device name (such as <code>Harddisk2</code>).
<code>p#c#t#l#</code>	The #s corresponds to the port, channel, target, and LUN of the disk.

Example

```
vxdisk -gDB1 merge Harddisk3
```

Merges Harddisk3 back into its former disk group, DB1.

vxdisk reactivate

```
vxdisk [-g<DynamicDiskGroupName>] reactivate <"Missing Disk (disk#)" |
DiskName |p#c#t#l#>
```

Reactivates a disk that has gone offline.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk reactivate` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	This parameter is required only if the internal disk name is used.
<code>"Missing Disk (disk#) "</code>	Specifies the missing disk. This string represents how the program identifies the missing disk. Because the string <code>"Missing Disk (disk#) "</code> has spaces, you must use quotation marks around it. The program specifies the (disk#), such as <code>(disk2)</code> .
<code><DiskName></code>	The internal disk name of the missing disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

Example

```
vxdisk reactivate "Missing Disk (disk2)"
```

Reactivates a disk that is shown as "Missing Disk (disk2)" that has gone offline. This assumes that the disk was offline and has now come back online.

vxdisk set hotreloc

```
vxdisk [-g<DynamicDiskGroupName>] set hotreloc
<mode=<enable|disable|on|off>>

| <usage=<preferred|secondary|nohotuse|reserved> <diskname>>
```

Set hot relocation at the host or disk level.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk set hotreloc` command.

Note: Dynamic disks belonging to a SAN Disk Group do not support the `vxdisk set hotreloc` command.

The following attributes apply:

<code>-g <DynamicDiskGroupName></code>	Specifies the dynamic disk group name.
<code>mode=<enable disable on off></code>	Specifies the mode to set the hot relocation. <ul style="list-style-type: none">■ Enable: Enable hot relocation at the host level.■ Disable: Disable hot relocation at the host level.■ On: Enable hot relocation at the disk level.■ Off: Disable hot relocation at the disk level.
<code>usage=<preferred secondary nohotuse reserved></code>	Hot relocation setting at the disk level. <ul style="list-style-type: none">■ Preferred: Preferred target disk in case of I/O failure.■ Secondary: Target disk that is used in case of I/O failure if no preferred disk is specified, or no free space available on preferred disks.■ Nohotuse: Does not allow disk to be used for hot relocation.■ Reserved: Does not allow disk to be used for hot relocation or other automatic selection operations.
<code><DiskName></code>	Specifies the disk name for hot relocation. (Requires the <code>-g</code> option; for example, <code>-gDG1 Disk2</code> .)

Example

```
vxdisk set hotreloc usage reserved disk2
```

Does not allow disk2 to be used for hot relocation.

vxdisk hotrelocinfo

```
vxdisk [-g<DynamicDiskGroupName>] hotrelocinfo [<diskname>]
```

The `vxdisk hotrelocinfo` CLI displays hot relocation settings for a given hard disk.

Example:

```
vxdisk hotrelocinfo harddisk1
```

Provides the information regarding `harddisk1` whether it is reserved for manual use or can be used as a hot relocation target. If the disk can be used as a hot relocation target, it provides information regarding its Hot Relocation usage.

The following attributes apply:

`-g<DynamicDiskGroupName>` Specifies the dynamic disk group name.

`<DiskName>` Specifies the disk name for hot relocation. (Requires the `-g` option; for example, `-gDG1 Disk2.`)

vxdisk sig

```
vxdisk sig "<DiskName> (no signature)
```

```
|p#c#t#l#> [diskstyle=<MBR|GPT>]
```

Adds a disk signature to a new disk. Because the string "`<DiskName> (no signature)`" has spaces, you must use quotation marks around it.

The following attributes apply:

`<DiskName>` The internal disk name of the disk.

`p#c#t#l#` Another way to indicate a disk; the `#s` correspond to port, channel, target, and LUN of a disk.

`diskstyle=<MBR|GPT>` Specifies the disk partitioning scheme. The default is MBR.

Example

```
vxdisk sig "Harddisk5 (no signature)"
```

Adds a signature to `Harddisk5` with a diskstyle of MBR.

vxdisk set track

```
vxdisk [-g<DynamicDiskGroupName>] set track align= <enable | disable  
| yes | no> offset=offset# <[vid=<vendor id> pid=<product id>] |  
<DiskName> | <p#c#t#l#>>
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Sets track alignment offset of data storage for disks that are connected to the host. When a disk is specified by name or by p#c#t#l#, the track alignment is set for all the disks that match the vendor ID and product ID of the disk. When the vendor ID and the product ID are specified, the track alignment is set for all the disks that match the specified vendor ID and product ID.

This command sets track alignment for all disks that are connected to the host if no disk or no vendor ID and product ID are specified.

The following attributes apply:

-g<DynamicDiskGroupName>	This parameter is required only if the internal disk name is used.
align	Enable or Disable track alignment. (Yes to enable or No to disable track alignment.)
vid	Vendor ID of storage device.
pid	Product ID of storage device.
offset	Track alignment offset of storage device (Range: 64 - 4096).
<DiskName>	Name of disk (for example, Harddisk2) or internal disk name (requires the -g option, for example, -gDG1 Disk2).
p#c#t#l#	Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

Examples

```
vxdisk set track align=enable vid=DGC pid=EMC offset=128
```

Enables track alignment and sets offset to 128 for all the disks that are connected to the host with vendor ID matching 'DGC' and product ID matching 'EMC'.

```
vxdisk set track align=enable vid=default pid=default offset=64
```

Enables track alignment and sets offset to 64 for all the disks that are connected to the host with vendor ID matching 'default' and product ID matching 'default'. (The 'default' specification refers to those disks that are not explicitly supported by the track alignment feature of SFW. Use the `vxdisk trackaligninfo` command to display which disks are supported by SFW.)

```
vxdisk set track align=enable offset=67 harddisk6
```

Enables track alignment and sets offset to 67 for all the disks that are connected to the host with vendor ID and product ID that match the vendor ID and product ID of `harddisk6`.

```
vxdisk -gdg1 set track align=enable offset=67 disk6
```

Enables track alignment and sets offset to 67 for all the disks that are connected to the host with vendor ID and product ID that match the vendor ID and product ID of `disk6` in diskgroup `dg1`.

```
vxdisk set track align=enable
```

Enables track alignment for all the disks that are connected to the host.

vxdisk trackaligninfo

```
vxdisk trackaligninfo
```

Displays the track alignment setting for the host and track alignment offset and settings for all arrays (including VID and PID).

Examples

```
vxdisk trackaligninfo
```

Displays the track alignment settings at the host level. Also displays the track alignment offset and track alignment setting for all arrays (identified by vendor ID and product ID).

vxdisk ping

```
vxdisk [-g<DynamicDiskGroupName>] ping <DiskName | p#c#t#l# >
```

Identifies a physical disk. This command flashes the fault light till a ping duration of 5 minutes is completed.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxdisk ping` command.

The following attributes apply:

<code>-g <DynamicDiskGroupName></code>	This parameter is required only if the internal disk name is used.
<code><DiskName></code>	The internal disk name of the disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

Example

```
vxdisk ping -gDG1 Disk2
```

Identifies Disk2 that belongs to disk group DG1.

vxdisk reclaim

```
vxdisk [-g<DynamicDiskGroupName>] [-b] reclaim<DiskName | p#c#t#l#>
```

Reclaims the storage space from a thin provisioned disk. The `-b` (background) option runs the reclaim operation in the background.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal disk name of the disk.)
<code>-b</code>	Run the command in the background.
<code><DiskName></code>	The DiskName can be the external name or the internal name. <ul style="list-style-type: none">■ The external name of the disk (for example, <code>Harddisk2</code>).■ The internal disk name of the disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).

p#c#t#l#

Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

Note: The operation to reclaim storage from thin provisioned disks is not supported for RAID-5, DCO, or Volume Replicator DCM volumes.

Note: The reclaim storage operation is disabled during a shrink volume operation. The reclaim storage operation is enabled when the shrink volume operation has completed.

Example

```
vxdisk reclaim Harddisk2 Harddisk3
```

This command reclaims space on the disks with the external names of Harddisk2 and Harddisk3.

vxdisk reclaimcancel

```
vxdisk [-g<DynamicDiskGroupName>] reclaimcancel<DiskName | p#c#t#l#>
```

Immediately cancels the reclaim operation. Only a portion of the storage space of the disk may have been reclaimed.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of dynamic disk group. (Needed only if you use the internal disk name of the disk.)
<code><DiskName></code>	The DiskName can be the external name or the internal name. <ul style="list-style-type: none"> ■ The external name of the disk (for example, Harddisk2). ■ The internal disk name of the disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).

p#c#t#l#

Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

Example

```
vxdisk reclaimcancel Harddisk2 Harddisk3
```

This command immediately cancels the reclaim operation on the disks with the external names of Harddisk2 and Harddisk3.

vxdisk setsite

To add an SFW dynamic disk to a site run the following command:

```
vxdisk [-g<DynamicDiskGroupName>] setsite site=<sitename> <diskname(s) | p#c#t#l# >
```

See [“Administering disks for sites through CLI”](#) on page 456.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Specifies the dynamic disk group name. For DG1
<code>site=<sitename></code>	Where site is the name of the site, for example, site1 . Sitename can take any alphanumeric values and can be maximum 31 characters long.
<code><diskname(s)></code>	The name of the disk or disks that is selected for tagging. For example, Harddisk1 , Harddisk2 , etc.
<code>p#c#t#l#</code>	Another way to indicate a disk; this corresponds to port, channel, target, and LUN of a disk.

Example:

```
vxdisk site1 Harddisk1 Harddisk2
```

This command adds disks Hard disk 1,Hard disk2 to site1.

vxdisk rmsite

Use `vxdisk rmsite` to remove a disk from the site.

Note: The `vxdisk rmsite` command fails on a disk that has Site Confined or Site Separated volumes that are associated with it.

```
vxdisk [-g<DynamicDiskGroupName>] rmsite <diskname(s) | p#c#t#l#>
```

See [“Administering disks for sites through CLI”](#) on page 456.

The following attributes apply:

<code><DynamicDiskGroupName></code>	This parameter is required only if the internal disk name is used.
<code><DiskName></code>	The internal disk name of the missing disk, which requires the <code>-g</code> option (for example, <code>-gDG1 Disk2</code>).
<code>rmsite</code> <code><diskname(s)></code>	Specifies the remove site or untag site option for disk names entered. For example, <code>rmsite Disk2</code> .
<code>p#c#t#l#</code>	Another way to indicate a disk; this corresponds to port, channel, target, and LUN of a disk.

Example:

```
vxdisk -gDG1 rmsite Disk2
```

This command removes Disk2 from its site.

vxassist

The `vxassist` utility has keywords or operands for creating and changing a volume and for doing operations that are related to a volume.

<code>make</code>	Creates a dynamic volume.
<code>growby</code>	Extends an NTFS or ReFS dynamic volume or an unformatted dynamic volume.
<code>querymax</code>	Queries the maximum number of reclaimable bytes from the current volume.
<code>shrinkby</code>	Decreases the size of an NTFS dynamic volume or an unformatted dynamic volume.
<code>mirror</code>	Adds a mirror to an existing volume.
<code>break</code>	Breaks a mirror from an existing volume.
<code>remove</code>	Removes a volume, a mirror, or a log.
<code>delete</code>	Deletes a partition or volume.
<code>shred</code>	Shreds and deletes a volume.
<code>addlog</code>	Adds a log plex to a volume.
<code>online</code> <code>(read/write)</code>	Sets the read/write access (Online) to a volume in a cluster-shared disk group (CSDG).

<code>offline</code> (<code>read-only</code>)	Sets the read-only access (Offline) to a volume in a cluster-shared disk group (CSDG).
<code>prepare</code>	Prepares a volume for a snapshot by adding a mirror to the volume. Note: Either the <code>prepare</code> or <code>snapstart</code> keyword may be used in the CLI, however <code>prepare</code> is recommended.
<code>snapshot</code>	Creates a separate volume from the snapshot mirror.
<code>snapback</code>	Joins the snapshot mirror back to the original volume.
<code>snapclear</code>	Clears the association between the original volume and the snapshot volume, making them separate volumes.
<code>snapabort</code>	Aborts the operation to prepare the snapshot on the original volume.
<code>rescan</code>	Rescans all the storage objects on the computer.
<code>refresh</code>	Refreshes the volume, drive letter, and system information.
<code>resetbus</code>	Performs a SCSI bus reset on all SCSI buses.
<code>version</code>	Displays the installed version of the product and the build number.

Windows-Specific Commands

<code>assign</code>	Assigns or changes a drive letter or drive path.
<code>unassign</code>	Removes a drive letter or drive path.
<code>create_part</code>	Creates a primary partition.
<code>create_ext</code>	Creates an extended partition.
<code>create_drive</code>	Creates a logical drive.
<code>delete_ext</code>	Deletes an extended partition.

Typing the following sequence for each keyword brings up a description of its syntax.

```
vxassist <Keyword> -?
```

vxassist make

Use the `vxassist make` command to create a volume.

```
vxassist [-b] [-o notrackalign] -g<DynamicDiskGroupName> make  
[<VolumeName>] <length>[type={mirror|raid5|stripe}[,log]]
```

```
[<alloc_attributes...>] [Column=N] [StripeUnit=<width>]
[Site=<sitename1>[,<sitename2>,...]]
[DriveLetter={A|B|..|Z}|DrivePath=<DrivePath>] [<[!]diskname |
p#c#t#l#> ...]
```

The following attributes apply:

<code>-b</code>	Option to run the command in the background. This option is useful for the operations that take an extended amount of time, such as adding a mirror or extending a volume.
<code>-o notrackalign</code>	Disable track alignment on the disk where the volume resides.
<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group.
<code><VolumeName></code>	Name of the volume. This is a Storage Foundation-specific name that is used for internal processing. The name is limited to 18 ASCII characters. It cannot contain spaces, slash mark (/), backslash (\), exclamation point (!), angle brackets (< >), or equal sign (=). Also, a period cannot be the first character in the name.
<code><Length></code>	Size of the volume; megabytes is the default. To indicate KB, add K to the number (500K); similarly, add G for GB and T for TB.
<code>type={mirror raid5 stripe} [,log]</code>	Type of volume to be created. Default is a spanned volume. The [,log] argument in type adds a RAID-5 log for a RAID-5 volume or a dirty region logging (DRL) log for a mirrored volume during volume creation.
<code><alloc_attributes></code>	Allocation attributes for mirroring and striping volumes. Attributes are: <ul style="list-style-type: none"> ■ Mirror=n[, {diskclass {[!]diskclass:instance}}] ■ Wantmirror={diskclass {[!]diskclass:instance}} ■ Stripe={diskclass {[!]diskclass:instance}} ■ Wantstripe={diskclass {[!]diskclass:instance}}

Mirror=<n>

or

Mirror=<n>, diskclass

Number of mirrors to be created for mirrored volumes. (Default is 2.)

Optionally specify a diskclass to use in mirroring. For example, `Mirror=2, target` specifies that volumes are mirrored between SCSI target addresses. Each mirror can contain disks from any number of instances of the diskclass, but different mirrors cannot use disks from the same instance of the diskclass.

Stripe=diskclass

Volumes are striped across instances of a particular diskclass. For example, `Stripe=target` specifies that volumes are striped between SCSI target addresses. Each column stripe can contain disks from any number of instances of the diskclass, but different columns cannot use disks from the same instance of the diskclass.

Wantmirror=diskclass

Specifies a diskclass to use in mirroring. However this constraint is ignored if the requested allocation cannot be satisfied.

Wantstripe=diskclass

Volumes are striped across instances of a particular disk class. However this constraint is ignored if the requested allocation cannot be satisfied.

```
diskclass  
{[!]diskclass:instance}
```

Diskclass specifies (or excludes) a particular type of disk grouping. For example, `port` specifies a group of disks on a port.

Instance specifies which grouping. For example, `p1` specifies a particular grouping of disks.

Each type of diskclass has a particular format for specifying instances of the class. Multiple diskclass instances can be specified separated by a space.

Diskclass can be one of the following:

- `target` or `t`

Specifies the disks that have the same SCSI target address of the same port. The target is specified in the form `p##t#`

- `enclr`, `e`, or `enclosure`

Specifies the disks belonging to a particular enclosure.

- `port` or `p`

Specifies the disks that are connected to a particular port.

- `channel` or `ch`

Specifies the disks that are connected to a particular channel.

```
Column=<n>
```

Number of columns. Required for RAID-5 and striped volumes.

```
StripeUnit=<Width>
```

Stripe width of RAID-5/striped volumes in blocks (512 Kbytes). Allows the user to specify the stripe unit for RAID-5 and striped volumes.

```
Site
```

Name of the site on which a site-based volume is to be created or mirrored.

```
DriveLetter={A|B|...|Z}
```

The default is no assignment of a drive letter to the volume.

```
DrivePath=<DrivePath>
```

Used when you want to mount the new volume at a folder on a local NTFS volume. You can use this attribute as an alternative to the drive letter.

[!] <DiskName>...

Name of a disk or disks to include or exclude, such as Harddisk2. Parameters are used to specify disks for volume creation, mirroring, and extension.

The disk name can also be indicated by the internal disk name or by p#c#t#l#, where the #s correspond to port, channel, target, and LUN of a disk.

If disks are not specified, SFW selects the disks automatically within the specified disk group.

Note: This attribute is required when working with a Microsoft Disk Management Disk Group.

Note: The default volume type is a concatenated volume. The created volume is not formatted (no file system is created on the volume). You need to use the operating system format command to format the volume.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Examples

```
vxassist -gDG2 make 100 DriveLetter=E Harddisk1
```

The above command line sequence creates a simple 100-MB volume in dynamic disk group DG2. The volume's drive letter is E. The volume would be created on Harddisk1.

```
vxassist -gDG2 make 100 type=mirror,log DriveLetter=F
```

The above command line sequence makes a mirrored 100-MB volume with a DRL log. The volume is assigned the drive letter F.

```
vxassist -gDG2 make 120 type=raid5 DriveLetter=X Column=3  
Harddisk2 Harddisk3 Harddisk5
```

The above command line sequence creates a 120-MB RAID-5 volume with the drive letter X. The volume is striped across Harddisk2, Harddisk3, and Harddisk5. All RAID-5 volumes require at least three columns. If you do not indicate a column number for a RAID-5 volume, the command fails.

```
vxassist -b -gDG2 make 100 type=stripe DriveLetter=X Column=3
```

The above command line sequence makes a striped volume that is assigned drive letter X. The `-b` option allows the command to run in the background.

```
vxassist -gDG2 make DbVolume 100
```

The resulting path name for the volume is:

```
\Device\HarddiskDmVolumes\DG2\DbVolume.
```

```
vxassist -gDG2 make Vol1 100m mirror=2, port
```

The above command line sequence makes a two mirrored 100-MB volumes that are connected to the same port.

```
vxassist -gDG2 make Vol1 100m mirror=2, port port:p2 port:p3
```

The above command line sequence makes two mirrored 100-MB volumes that are connected to port:p2 and port:p3.

To create a volume for a site that uses the site-aware allocation feature, refer to:

See [“Creating a volume through CLI”](#) on page 459.

See [“Administering volumes for site-based allocation ”](#) on page 456.

vxassist growby

```
vxassist [-b] [-o notrackalign] [-g<DynamicDiskGroupName>] growby
<VolumeName | DriveLetter | VmName|DrivePath> <Length>
[<alloc_attributes...>] [[!]<DiskName|p#c#t#l#>]...
```

Extends the usable space of an NTFS, an ReFS, or an unformatted volume by the amount that is specified in *<Length>*.

The following attributes apply:

<code>-b</code>	Runs the command in the background.
<code>-o notrackalign</code>	Disable track alignment on the disk where the volume resides.
<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<i><VolumeName></i>	the path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<i><DriveLetter></i>	The drive letter of the volume.

<VmName>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<DrivePath>	A volume that is mounted on an NTFS or ReFS folder; it is indicated by a path to that folder.
<Length>	The size of the volume; megabytes is the default. To indicate KB, add K to the number (500K); similarly, add G for GB and T for TB.
<alloc_attributes>	<p>Allocation attributes for mirroring and striping volumes.</p> <p>Attributes are:</p> <ul style="list-style-type: none"> ■ Mirror={diskclass {[!]diskclass:instance}} ■ Wantmirror={diskclass {[!]diskclass:instance}} ■ Stripe={diskclass {[!]diskclass:instance}} ■ Wantstripe={diskclass {[!]diskclass:instance}}
Mirror=diskclass	Specifies a diskclass to use in mirroring. For example, <code>Mirror=target</code> specifies that volumes are mirrored between SCSI target addresses. Each mirror can contain disks from any number of instances of the diskclass, but different mirrors cannot use disks from the same instance of the diskclass.
Stripe=diskclass	Volumes are striped across instances of a particular diskclass. For example, <code>Stripe=target</code> specifies that volumes are striped between SCSI target addresses. Each column stripe can contain disks from any number of instances of the diskclass, but different columns cannot use disks from the same instance of the diskclass.
Wantmirror=diskclass	Specifies a diskclass to use in mirroring. However this constraint is ignored if the requested allocation cannot be satisfied.
Wantstripe=diskclass	Volumes are striped across instances of a particular disk class. However this constraint is ignored if the requested allocation cannot be satisfied.

```
diskclass  
{[!]diskclass:instance}
```

Diskclass specifies (or excludes) a particular type of disk grouping. For example, `port` specifies a group of disks on a port.

Instance specifies which grouping. For example, `p1` specifies a particular grouping of disks.

Each type of diskclass has a particular format for specifying instances of the class. Multiple diskclass instances can be specified separated by a space.

Diskclass can be one of the following:

- `target` or `t`

Specifies the disks that have the same SCSI target address of the same port. The target is specified in the form `p##`

- `enclr`, `e`, or `enclosure`

Specifies the disks belonging to a particular enclosure.

- `port` or `p`

Specifies the disks that are connected to a particular port.

- `channel` or `ch`

Specifies the disks that are connected to a particular channel.

```
Column=<n>
```

Number of columns. Required for RAID-5 and striped volumes.

```
StripeUnit=<Width>
```

Stripe width of RAID-5/striped volumes in blocks (512 Kbytes). Allows the user to specify the stripe unit for RAID-5 and striped volumes.

```
[!]<DiskName>...
```

The name of the disk or disks to include or exclude on which the volume is extended. The disk name can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name the (for example, `-gDG1 Disk2`).

```
[!]<p#c#t#l#>
```

Another way to indicate a disk to include or exclude; the `#s` correspond to port, channel, target, and LUN of the disk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Examples

```
vxassist growby J: 100
```

The above command extends volume J by 100 MB. Only NTFS, ReFS, or unformatted volumes can be extended. An unformatted volume's file system type is shown as "RAW."

```
vxassist -g mydg growby vol1 100M port:p3 !Harddisk2
```

The above command extends vol1 by 100 MB on disks that are connected to port p3 excluding Harddisk2.

vxassist querymax

```
vxassist [-g<DynamicDiskGroupName>] querymax <VolumeName | DriveLetter  
| VmName | DrivePath>
```

Queries the maximum number of reclaimable bytes from the current volume. This command supports only NTFS, ReFS, or unformatted volumes. An unformatted volume's file system type is shown as "RAW."

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<code><VolumeName></code>	the path name of the volume, such as <code>\Device\HarddiskDmVolumes\DGL\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDGL Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS or ReFS folder; it is indicated by a path to that folder.

Examples

```
vxassist querymax J:
```

The above command queries volume J maximum number of reclaimable bytes. The output is the number of reclaimable bytes that can be reclaimed.

The output of `vxassist querymax` may not be accurate. The command internally uses the Microsoft API to query the volumes. These APIs can sometimes return more reclaimable bytes than the actual available bytes.

See www.support.microsoft.com/kb/976436

vxassist shrinkby

```
vxassist [-g<DynamicDiskGroupName>] shrinkby <VolumeName | DriveLetter  
| VmName|DrivePath> [<Length>]
```

Decreases an NTFS or unformatted volume (RAW) by the amount that is specified in `<Length>`. If the optional parameter `<Length>` is not used, then the command decreases the volume by maximum reclaimable bytes. Before using the command, review the prerequisites for shrinking a dynamic volume.

See “[Before shrinking a dynamic volume](#)” on page 131.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<code><VolumeName></code>	the path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

<Length>

The size by which you want to shrink the volume. Acceptable units are "S" for Sectors, "K" for KBs, "M" for MBs, "G" for GBs, "T" for TBs. Default unit is MB. For example, to reduce a 30 GB volume to 20 GB, provide 10G.

If not specified, then `vxassist querymax` is used to obtain maximum reclaimable bytes for a given volume.

`vxassist querymax` internally uses Microsoft APIs to query a volume. These APIs can sometimes return more reclaimable bytes than the actual available bytes.

See www.support.microsoft.com/kb/976436

In such a case the `shrinkby` operation may fail with the following error: V-76-42261-7777: Shrink volume Failed: No room to shrink.

To resolve this issue, run the `shrinkby` operation again or specify a size that is lower than the value that is obtained from the `vxassist querymax` command.

Note: The shrink volume operation is disabled during a reclaim storage operation for thin provisioned disks. The shrink volume operation is enabled when the reclaim storage operation has completed.

Note: During the volume shrink operation, if the deport disk group operation is performed, either by the user or because of a failover, then the volume shrink operation is canceled.

Examples

```
vxassist shrinkby J: 100
```

The above command shrinks volume J by 100 MB. This command supports only NTFS or unformatted volumes. An unformatted volume's file system type is shown as "RAW."

vxassist shrinkabort

```
vxassist [-g<DynamicDiskGroupName>] shrinkabort
```

<VolumeName|DriveLetter|VmName|DrivePath>

The `vxassist shrinkabort` cancels the shrink operation on the given volume.

The following attributes apply:

`DynamicDiskGroupName` Name of the dynamic disk group.

`VolumeName` Name of the volume.

Example

```
\Device\HarddiskDmVolumes\DynamicGroup\Volume1
```

`DriveLetter` Drive letter of the volume.

Example

```
C:
```

`VmName` Internal name of volume.

(Requires use of `-g<DynamicDiskGroupName>` option),

Example

```
-g MyDg Volume1
```

`DrivePath` Drive Path on which the volume is mounted on.

```
C:\TEMP
```

Must be an empty folder on a local NTFS volume.

Example

```
vxassist shrinkabort z:
```

The `shrink` volume operation is performed in two stages, where the `filesystem shrink` operation is performed after which `volumeshrink` follows.

If the `filesystem shrink` operation has progressed too far to cancel, then it commits the changes that are made during the `filesystem shrink` operation after which `volumeshrink` operation is aborted.

In such a case, use `vxvol growfs` to extend file system to volume size.

vxassist mirror

Use the `vxassist mirror` command to add a mirror to an existing volume.

```
vxassist [-b] [-o notrackalign] [-g<DynamicDiskGroupName>] mirror  
<VolumeName|DriveLetter|VmName|DrivePath> [type=<typename>]
```

```
[<alloc_attributes ...>] [Column=N] [StripeUnit=<width>]
[Site=<sitename1>[,<sitename2>,...]] [<[!]diskname | p#c#t#l#> ...]
```

The following attributes apply:

-b	Runs the command in the background.
-o notrackalign	Disable track alignment on the disk where the volume resides.
-g<DynamicDiskGroupName>	Required only if the internal volume name or internal disk name is used.
<VolumeName>	the path name of the volume, such as \Device\HarddiskDmVolumes\DG1\Volume1.
<DriveLetter>	The drive letter of the volume.
<VmName>	Internal name of the volume; requires the use of the -g option (for example, -gDG1 Volume1).
<DrivePath>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
Type=<Type>	Type of volume layout. Types are: - Stripe - Concatenate
<alloc_attributes>	Allocation attributes for mirroring and striping volumes. Attributes are: <ul style="list-style-type: none"> ■ Mirror=n[, {diskclass {[!]diskclass:instance}}] ■ Wantmirror={diskclass {[!]diskclass:instance}} ■ Stripe={diskclass {[!]diskclass:instance}} ■ Wantstripe={diskclass {[!]diskclass:instance}}

`Mirror=<n>`

or

`Mirror=<n>, diskclass`

Number of mirrors to be created for mirrored volumes. (Default is 2.)

Optionally specify a diskclass to use in mirroring. For example, `Mirror=2, target` specifies that volumes are mirrored between SCSI target addresses. Each mirror can contain disks from any number of instances of the diskclass, but different mirrors cannot use disks from the same instance of the diskclass.

`Stripe=diskclass`

Volumes are striped across instances of a particular diskclass. For example, `Stripe=target` specifies that volumes are striped between SCSI target addresses. Each column stripe can contain disks from any number of instances of the diskclass, but different columns cannot use disks from the same instance of the diskclass.

`Wantmirror=diskclass`

Specifies a diskclass to use in mirroring. However this constraint is ignored if the requested allocation cannot be satisfied.

`Wantstripe=diskclass`

Volumes are striped across instances of a particular disk class. However this constraint is ignored if the requested allocation cannot be satisfied.


```
diskclass
{[!]diskclass:instance}
```

Diskclass specifies (or excludes) a particular type of disk grouping. For example, `port` specifies a group of disks on a port.

Instance specifies which grouping. For example, `p1` specifies a particular grouping of disks.

Each type of diskclass has a particular format for specifying instances of the class. Multiple diskclass instances can be specified separated by a space.

Diskclass can be one of the following:

- `target` or `t`

Specifies the disks that have the same SCSI target address of the same port. The target is specified in the form `p###`

- `enclr`, `e`, or `enclosure`

Specifies the disks belonging to a particular enclosure.

- `port` or `p`

Specifies the disks that are connected to a particular port.

- `channel` or `ch`

Specifies the disks that are connected to a particular channel.

```
Column=<n>
```

Number of columns. Required for RAID-5 and striped volumes.

```
StripeUnit=<Width>
```

Stripe width of RAID-5/striped volumes in blocks (512 Kbytes). Allows the user to specify the stripe unit for RAID-5 and striped volumes.

```
Site
```

Specifies the name of the site on which a site-based volume is to be created or mirrored.

[!] <DiskName>...

The name of the disk or disks (such as `Harddisk2`) on which the mirror or mirrors are to be created.

The disk name can also be indicated by the internal disk name, which requires the `-g` option (for example, `-gDG1 Disk2`).

[!] <p#c#t#l#>

Another way to indicate a disk; the `#s` correspond to port, channel, target, and LUN of a disk.

Warning: Adding mirrors to a volume involves a certain amount of time for mirror resynchronization. Take care not to perform actions on the volume until the command is complete. Also, if you inadvertently shut down the server or deport the disk group containing the volume before mirror resynchronization completes, any mirrors that were being added to the volume are deleted when the disk group with the volume comes online again.

The SmartMove feature can help reduce mirror resynchronization time.

See [“SmartMove”](#) on page 93.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support track alignment.

Example

```
vxassist mirror -b Z: Mirror=2 Harddisk4 Harddisk5
```

The above command line sequence adds two mirrors to volume Z. `Harddisk4` and `Harddisk5` are used to create the mirrors. The process of adding the two mirrors runs in the background. Instead of using the drive letter Z, you can provide a path name, such as:

```
\Device\HarddiskDmVolumes\DynDskGrp1\Volume1
```

To create a mirror for a site-based volume using the site-aware allocation feature, refer to:

See [“Adding a mirror to a site-based volume through CLI”](#) on page 465.

See [“Administering mirrored plexes for site-based allocation”](#) on page 460.

vxassist break

```
vxassist [-g<DynamicDiskGroupName>] break <VolumeName | DriveLetter>
<plex=PlexName>|<!Diskname |!p#c#t#l#>]
```

```
[DriveLetter={A|B|...|Z}|DrivePath=<DrivePath>]
```

Breaks a mirror from an existing volume. When a mirror is broken, it becomes a separate volume and you lose the fault tolerance that mirroring provides.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code>plex=<PlexName></code>	The mirror or plex to be broken off. It can be the plex name (such as <code>Volume2-01</code>) or the GUID of the mirror plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command <code>vxvol volinfo</code> for the mirrored volume that contains the plex. If you use a script and have multiple snapshot plexes, then the GUID can distinguish the different snapshot plexes.
<code><!DiskName></code>	The name of the disk, such as <code>Harddisk2</code> , or the internal disk name, which requires the use of the <code>-g</code> option (for example, <code>!-gDG1 Disk2</code>). You can use this attribute as an alternative to the plex name. See the note on the exclamation mark that follows this table.
<code><!p#c#t#l#></code>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of the disk. It also requires the <code>!</code> .
<code>DriveLetter={A B ... Z}</code>	Used to assign a specific drive letter to the new volume that is created from the broken-off plex. The drive letter is not automatically assigned.

DrivePath=<DrivePath>

Used when you want to mount the new volume that is created from the broken-off plex at a folder on a local NTFS volume. You can use this attribute as an alternative to the drive letter.

Note: The exclamation mark in front of the *DiskName* attribute is necessary to make the command compatible with the Volume Manager for UNIX version of the command.

If you do not remember the plex name, you can determine what it is by using the `vxvol volinfo` command. Alternatively, if you know which disk the plex resides on, you can use the `<!DiskName>` syntax, for example, `!Harddisk2`, to specify the plex.

Note: If you want to delete the mirror instead of breaking it off, see the `vxremove` command in the section.

Examples

```
vxassist break H: plex=Volume1-01 DriveLetter=Z
```

The above command line sequence breaks a mirror from a volume that has been assigned the drive letter H. The specified plex becomes a new volume and receives the drive letter Z. The original volume retains the drive letter H.

```
vxassist break\Device\HarddiskDmVolumes\DynDskGrp1\Volume1  
plex=Volume1-01 DriveLetter=Z
```

The above command line sequence is the same as the previous command sequence except that a path name is used to indicate the volume from which the mirror is broken.

```
vxassist break L: !Harddisk3 DriveLetter=M
```

This breaks off the mirror that resides on Harddisk3 from volume L and assign drive letter M to the newly created volume.

vxassist remove

```
vxassist [-f] [-g<DynamicDiskGroupName>] remove [volume|mirror|log]  
<VolumeName | DriveLetter | VmName|DrivePath> [LogType=<DRL|DCM |DCO>]  
[nlog=<#>] [plex=<PlexName>|<!Diskname |!p#c#t#l#>]
```

Removes (deletes) either a volume, a mirror from an existing volume, or a DRL or RAID-5 log from an existing volume. When you specify removing a volume, the command works the same as `vxassist delete`. If no volume, mirror, or log is specified, the command assumes a mirror plex for backward compatibility because the command formerly removed only a mirror plex.

The following attributes apply:

<code>-f</code>	Force the operation.
<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<code>volume mirror log</code>	Specifies what is to be removed from an existing volume. If you do not use this parameter, a mirror is removed. If you want to remove a volume or a log, you have to put the <code>volume</code> or <code>log</code> keyword in.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume to be operated on.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code>LogType</code>	Specify the type of log when removing a log (DCM, DRL, or DCO type).
<code>nlog=<n></code>	Specify the number of logs when removing a log. Number of logs. (n>0)
<code>plex=<PlexName></code>	The mirror or plex to be removed. It can be the plex name (such as <code>Volume2-01</code>) or the GUID of the mirror plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command <code>vxvol volinfo</code> for the mirrored volume that contains the plex. If you use a script and have multiple snapshot plexes, then the GUID can distinguish the different snapshot plexes.

<code><!DiskName></code>	The name of the disk, such as <code>Harddisk2</code> , or the internal disk name, which requires the use of the <code>-g</code> option (for example, <code>!-gDG1 Disk2</code>). You can use this attribute as an alternative to the plex name. See the note on the exclamation mark that follows this table. If there is more than one mirror or log associated with the volume, you need to indicate either the mirror plex name or log plex name or the disk name.
<code><!p#c#t#l#></code>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of the disk. It also requires the <code>!</code> .

Note: The exclamation mark in front of the `DiskName` attribute is necessary to make the command compatible with the Volume Manager for UNIX version of the command.

If you do not remember the plex name, you can determine what it is by using the `vxvol volinfo` command. Alternatively, if you know which disk the plex resides on, you can use the `<!DiskName>` syntax, for example, `!Harddisk2`, to specify the plex.

The logs can be either DRL or RAID-5 logs. If you have more than one log, you need to identify the log plex. You cannot remove a DCO log with this command. A DCO log is added with the command `vxvol set fastresync=on` and deleted with the command `vxvol set fastresync=off`.

See “[vxvol](#)” on page 581.

Examples

```
vxassist remove Z: plex=Volume1-01
```

The above command line sequence removes a mirror from volume Z. The data on the specified plex is destroyed and is not recoverable.

```
vxassist remove volume F:
```

The above command line sequence removes volume F. The data on the volume is destroyed and is not recoverable.

```
vxassist remove log H:
```

The above command line sequence removes a log from volume H.

vxassist delete

```
vxassist [-f] [-g<DynamicDiskGroupName>] delete  
<VolumeName | DriveLetter | VmName|DrivePath>
```

Deletes a volume or partition. The following attributes apply:

<code>-f</code>	Force the operation.
<code>-g<DynamicDiskGroupName></code>	Required only if the <code><VmName></code> parameter is used.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume to be operated on.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Examples

```
vxassist delete J:
```

The above command deletes volume J. The command destroys all data on the volume.

```
vxassist delete\Device\HarddiskDmVolumes\DynDskGrp1\volume1
```

The above command deletes a volume with the path name:

```
\Device\HarddiskDmVolumes\DynDskGrp1\volume1.
```

vxassist shred

```
vxassist [-f] [-g<DynamicDiskGroupName>] shred [shredtype=<0|1|2>]  
<VolumeName | DriveLetter | VmName|DrivePath>
```

Overwrites the entire volume with zeroes and then deletes the volume.

SFW provides the following methods for overwriting the volume:

- 1 Pass algorithm
Overwrites the volume with a pattern of zeroes.

Specified by `shredtype=0`.

- 3 Pass algorithm

The volume is overwritten a total of 3 times. The volume is first overwritten with a pre-selected digital pattern, then overwritten with the binary complement of the pattern. In the last pass, the volume is overwritten with a randomly selected digital pattern. This 3 pass approach is based on the specifications in the US Department of Defense standard 5220.22-M (3).

Specified by `shredtype=1`.

- 7 Pass algorithm

The volume is overwritten a total of 7 times. Each pass consists of overwriting the volume with a randomly selected digital pattern or with the binary complement of the previous pattern. This 7 pass approach is based on the specifications in the US Department of Defense standard 5200.28-STD (7).

Specified by `shredtype=2`.

SFW defaults to the 1 Pass algorithm when `shredtype` is not specified.

Warning: All data in the volume is 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.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist shred` command.

The following attributes apply:

<code>-f</code>	Force the operation.
<code>-g<DynamicDiskGroupName></code>	Required only if the <code><VmName></code> parameter is used.
<code>-shredtype=<0 1 2></code>	Specifies the method to overwrite the volume. <ul style="list-style-type: none"> ■ 0 specifies the 1 Pass algorithm ■ 1 specifies the 3 Pass algorithm ■ 2 specifies the 7 Pass algorithm
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume to be operated on.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).

<DrivePath> A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Examples

```
vxassist shred J:
```

The above command shreds volume J.

```
vxassist shred\Device\HarddiskDmVolumes\DynDskGrp1\volume1
```

The above command shreds a volume with the path name:

```
\Device\HarddiskDmVolumes\DynDskGrp1\volume1.
```

vxassist addlog

```
vxassist [-o notrackalign] [-g<DynamicDiskGroupName>] addlog
<VolumeName | DriveLetter | VmName|DrivePath> [LogType=<DRL | DCM |
DCO>] [nlog=<n>] [[!]diskclass:instance] [[!]<DiskName|p#c#t#l#>...]
```

Adds a DRL, DCM, DCO, or RAID-5 log plex to the volume.

RAID-5 logging is used to prevent bad volumes during system crash or failure. If more than one log is added, then the logs are just mirrored. DRL logs for RAID volumes are called RAID-5 logs and provide the same functionality.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist addlog` command.

The following attributes apply:

<code>-o notrackalign</code>	Disable track alignment on the disk where the log resides.
<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or internal disk name is used.
<VolumeName>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<DriveLetter>	The drive letter of the volume to be operated on.
<VmName>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<DrivePath>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

LogType	Type of log to add (DCM, DRL, or DCO type).
nlog=<n>	Number of logs to add.
[!]diskclass:instance	Specify (or exclude) a set of disks. Diskclass specifies a particular type of disk grouping. For example, <code>port</code> specifies a group of disks on a port. Instance specifies which grouping. For example, <code>p1</code> specifies a particular grouping of disks. Each type of diskclass has a particular format for specifying instances of the class. Multiple diskclass instances can be specified separated by a space.
diskclass	Diskclass can be one of the following: - <code>target</code> or <code>t</code> Specifies the disks that have the same SCSI target address of the same port. The target is specified in the form <code>p##t##</code> - <code>enclr</code> , <code>e</code> , or <code>enclosure</code> Specifies the disks belonging to a particular enclosure. - <code>port</code> or <code>p</code> Specifies the disks that are connected to a particular port. - <code>channel</code> or <code>ch</code> Specifies the disks that are connected to a particular channel.
[!]<DiskName>...	Specifies the disk or disks to include or exclude. If you do not use this parameter, the program selects the disks. The <i>DiskName</i> parameter can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG1 Disk2</code>).
[!]<p#c#t#l#>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of a disk.

For descriptions of the advanced features of DRL and RAID-5 logging, see the following topic:

See [“Dirty region logging \(DRL\) and RAID-5 logging”](#) on page 222.

Example

```
vxassist -g mydg addlog voll logtype=DRL port:p2 !harddisk3
```

Adds a DRL log to vol1 on a disk that is not harddisk3 and is connected to port p2.

vxassist online (read/write)

```
vxassist [-g<ClusterSharedDiskGroupName>]
```

```
online <VolumeName | DriveLetter | VmName | DrivePath>
```

Brings a volume in a cluster-shared disk group (CSDG) online and sets read/write access to it. The following attributes apply:

-g<ClusterSharedDiskGroupName> Name of cluster-shared disk group. (Needed only if you use the internal volume name parameter.)

<VolumeName> The path name of the volume, such as
\\Device\HarddiskDmVolumes\DG1\Volume1.

<DriveLetter> The drive letter of the volume.

<VmName> Internal name of the volume; requires the use of the
-g option, such as -gCSDG1 Volume1.

<DrivePath> A volume that is mounted on an NTFS folder; it is
indicated by a path to that folder.

Example:

```
vxassist -gCSDG1 online Volume1
```

This command sets read/write access to the volume "Volume1" of the cluster-shared disk group "CSDG1" and brings it online.

vxassist offline

```
vxassist [-f] [-g<ClusterSharedDiskGroupName>]
```

```
offline <VolumeName | DriveLetter | VmName | DrivePath>
```

Brings a volume in a cluster-shared disk group (CSDG) offline. The following attributes apply:

-g<ClusterSharedDiskGroupName> Name of cluster-shared disk group. (Needed only if you use the internal volume name parameter.)

<code>[-f]</code>	The force option, forces the volume to be brought offline.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DG1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gCSDG1 Volume1.</code>
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Example:

```
vxassist -gCSDG1 offline Volume1
```

This command brings offline the volume "Volume1" of the cluster-shared disk group "CSDG1".

vxassist prepare

```
vxassist [-b] [-o notrackalign] [-g<DynamicDiskGroupName>] prepare
<VolumeName | DriveLetter | VmName|DrivePath>
[Type=stripe][<alloc_attributes...>] [Column=<n>]
[StripeUnit=<Width>][Plex=<MirrorPlexName> | [!]<Diskname> | p#c#t#l#>
...]
```

Adds a snapshot mirror to the volume. It is a mirrored plex of the original volume that is broken off in the snapshot process. You can either add a new snapshot mirror or convert a mirror of the volume that already exists. A snapshot mirror is also referred to as a "snap plex."

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist prepare` command.

Note: Either the `prepare` or `snapstart` keyword may be used in the CLI, however `prepare` is recommended.

The following attributes apply:

<code>-b</code>	Runs the command in the background.
-----------------	-------------------------------------

<code>-o notrackalign</code>	Disable track alignment on the disk where the mirror resides.
<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name or internal disk name.
<code><VolumeName></code>	The path name of the volume, such as <code>\Device\HarddiskDmVolumes\DynDskGrp1\Volume1.</code>
<code><DriveLetter></code>	The drive letter of the existing volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code>Type=<Type></code>	Type of volume layout. Types are: <ul style="list-style-type: none"> - Stripe - Spanned (Spanned is the default.)
<code><alloc_attributes></code>	Allocation attributes for mirroring and striping volumes. Attributes are: <ul style="list-style-type: none"> ■ <code>Mirror={diskclass {[!]diskclass:instance}}</code> ■ <code>Wantmirror={diskclass {[!]diskclass:instance}}</code> ■ <code>Stripe={diskclass {[!]diskclass:instance}}</code> ■ <code>Wantstripe={diskclass {[!]diskclass:instance}}</code>
<code>Mirror=diskclass</code>	Specifies a diskclass to use in mirroring. For example, <code>Mirror=target</code> specifies that volumes are mirrored between SCSI target addresses. Each mirror can contain disks from any number of instances of the diskclass, but different mirrors cannot use disks from the same instance of the diskclass.
<code>Stripe=diskclass</code>	Volumes are striped across instances of a particular diskclass. For example, <code>Stripe=target</code> specifies that volumes are striped between SCSI target addresses. Each column stripe can contain disks from any number of instances of the diskclass, but different columns cannot use disks from the same instance of the diskclass.
<code>Wantmirror=diskclass</code>	Specifies a diskclass to use in mirroring. However this constraint is ignored if the requested allocation cannot be satisfied.

<code>Wantstripe=diskclass</code>	Volumes are striped across instances of a particular disk class. However this constraint is ignored if the requested allocation cannot be satisfied.
<code>diskclass</code> {[!]diskclass:instance}	<p>Diskclass specifies (or excludes) a particular type of disk grouping. For example, <code>port</code> specifies a group of disks on a port.</p> <p>Instance specifies which grouping. For example, <code>p1</code> specifies a particular grouping of disks.</p> <p>Each type of diskclass has a particular format for specifying instances of the class. Multiple diskclass instances can be specified separated by a space.</p> <p>Diskclass can be one of the following:</p> <ul style="list-style-type: none"> - <code>target</code> or <code>t</code> <p>Specifies the disks that have the same SCSI target address of the same port. The target is specified in the form <code>p###</code></p> <ul style="list-style-type: none"> - <code>enclr</code>, <code>e</code>, or <code>enclosure</code> <p>Specifies the disks belonging to a particular enclosure.</p> <ul style="list-style-type: none"> - <code>port</code> or <code>p</code> <p>Specifies the disks that are connected to a particular port.</p> <ul style="list-style-type: none"> - <code>channel</code> or <code>ch</code> <p>Specifies the disks that are connected to a particular channel.</p>
<code>Column=<n></code>	Number of columns. Required for RAID-5 and striped volumes.
<code>StripeUnit=<Width></code>	Stripe width of RAID-5/striped volumes in blocks (512 Kbytes). Allows the user to specify the stripe unit for RAID-5 and striped volumes.
<code>plex=<MirrorPlexName></code>	Converts a specified mirror plex to a snap plex. It can be the plex name (such as <code>Volume1-01</code>) or the GUID of the mirror plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command <code>vxvol -v volinfo</code> for the mirrored volume that contains the plex.
<code>[!]<DiskName></code>	The designated hard disk to include or exclude, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG1 Disk2</code>).

[!]<p#c#t#l#> Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of a disk.

If you do not remember the plex name, you can determine what it is by using the `vxvol volinfo` command.

When the snapshot mirror is synchronized with the volume, its state changes to Snap Ready and the prepare the volume for snapshot task exits.

More information is available on Snap Shot commands.

See [“Snapshot commands”](#) on page 259.

Examples

```
vxassist -b prepare E:
```

Runs the prepare the volume for snapshot process on volume E in the background.

```
vxassist prepare E: plex=Volume2-01
```

Runs the prepare the volume for snapshot process on volume E, converting the existing mirror plex Volume to a snap plex.

```
vxassist prepare e: plex=volume1-02
```

This command converts the plex named "volume1-02" from a regular mirror to a snapshot plex.

```
vxassist prepare e: harddisk2 harddisk3
```

This command creates a new snap plex on harddisk2 and/or harddisk3. You can indicate more than one disk to be used for the snap plex.

```
vxassist prepare d:\mount1
```

This command adds a snap plex to the volume that is mounted on folder d:\mount1.

Note: Either the `prepare` or `snapstart` keyword may be used in the CLI, however `prepare` is recommended.

Note: You should not use the `vxassist break` command to break off a snapshot mirror from its associated volume or the `vxassist remove` command to delete a snapshot mirror. Instead use `vxassist snapshot` to break off a snapshot mirror from the associated volume or use `vxassist snapabort` to delete it.

vxassist snapshot

```
vxassist [-g<DynamicDiskGroupName>] snapshot
<VolumeName|DriveLetter|VmName|DrivePath...>
[plex=<SnapPlexName>|<!DiskName |!p#c#t#l# >] [access=readonly]
[DriveLetter={A|B|...|Z}]|DrivePath=<DrivePath>] [<NewVolumeName>]
```

Detaches the snapshot mirror (also called the "snap plex") and creates a separate volume that is associated with the snapshot mirror.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist snapshot` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name or internal disk name parameter.
<code><VolumeName></code>	The path name of the original volume, such as <code>\Device\HarddiskDmVolumes\DynDskGrp1\Volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code>plex=<SnapPlexName></code>	Name of the snapshot mirror. It can be the plex name (such as <code>Volume1-01</code>) or the GUID of the mirror plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command <code>vxvol -v volinfo</code> for the mirrored volume that contains the plex. You need to specify this attribute only when there is more than one snap plex.
<code><!DiskName></code>	The name of the disk, such as <code>Harddisk2</code> , or the internal disk name, which requires the use of the <code>-g</code> option (for example, <code>!-gDG1 Disk2</code>). See the note below this table on the exclamation mark. You can use <code>!DiskName</code> as an alternative to the plex name.

<code><!p#c#t#l#></code>	Another way to indicate a disk; the #s correspond to port, channel, target, and LUN of the disk. It also requires the !.
<code>access=readonly</code>	Sets the new volume to read only access.
<code>DriveLetter={A B ... Z}</code>	Used to assign a specific drive letter to the new volume.
<code>DrivePath=<DrivePath></code>	Used to assign a drive path to the new volume.
<code><NewVolumeName></code>	The name of the volume that is created when the snapshot mirror is broken off. This is also referred to as the <i>SnapVolumeName</i> parameter.

Examples

```
vxassist snapshot E: access=readonly DriveLetter=H
```

This command detaches a snapshot mirror from volume E. The snapshot volume is volume H with read only access.

```
vxassist snapshot E: plex=volume1-02 DrivePath=C:\ABC BackupVolume
```

This command detaches the snapshot mirror, volume1-02, from volume E, and mounts the new volume called "BackupVolume" to an empty folder called "ABC" on the C drive. The device path of the new volume is:

```
\Device\HarddiskDMVolumes\DynDskGrp1\BackupVolume.
```

```
vxassist snapshot d:\mp1 snap
```

This snapshots the volume which is mounted at the d:\mp1 folder and assigns the "snap" name to the snapshot.

vxassist snapback

```
vxassist [-o resyncfromreplica] [-b] [-f] [-g<DynamicDiskGroupName>]  
snapback <SnapVolumeName | DriveLetter | VmName | DrivePath>
```

Joins the broken plex back to the original volume and resynchronizes the two volumes. By default, the synchronization is done from the original volume to the snapshot mirror.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist snapback` command.

The following attributes apply:

<code>-o resyncfromreplica</code>	Reverses the resynchronization, copying data from the snapshot mirror to the original volume
<code>-b</code>	Performs the command in the background.
<code>-f</code>	Forces the snapback. Use this option with care. Make sure that the volume is not in use.
<code>-g<DynamicDiskGroupName></code>	Needed only if you use an internal volume name.
<code><SnapVolumeName></code>	The name of the volume that is associated with the snapshot plex.
<code><DriveLetter></code>	The drive letter of the snapshot volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

Example

```
vxassist snapback H:
```

This command joins the snapshot volume H back to the original volume.

```
vxassist snapback d:\mount1
```

This command snaps back the plex that is associated with the volume that is mounted on the d:\mount1 folder.

vxassist snapclear

```
vxassist -g<DynamicDiskGroupName> snapclear
```

```
<VolumeName | SnapVolumeName | DriveLetter | VmName | DrivePath>  
[plex=<SnapPlexName>] [<SnapshotName>]
```

Clears the association between the original volume and the snapshot volume and makes the snapshot volume and original volume two independent volumes. The `vxassist snapclear` command can be specified only after the `vxassist snapshot` command has been issued. After the `vxassist snapclear` command runs, the snapshot volume becomes an ordinary volume and cannot be snapped back to the original volume.

With SFW, it is possible to do a snapclear operation on the original volume. You might want to clear the original volume in two situations: 1) A snapshot volume is

deported to a second host and then deleted on that host, and 2) A snapclear is done on the snap volume but the original volume is offline and later this volume comes online.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist snapclear` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Needed if you use the internal volume name or original volume name.
<code><VolumeName></code>	The name of the original volume.
<code><SnapVolumeName></code>	The name of the volume that is associated with the snapshot plex, the volume that is broken off through the <code>vxassist snapshot</code> command.
<code><DriveLetter></code>	The drive letter of the snapshot volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code>plex=<SnapPlexName></code>	Name or GUID of the mirror plex.
<code><SnapshotName></code>	Optional name that is assigned to the snapshot volume.

Examples

```
vxassist snapclear H:
```

This command clears the association between the snapshot volume H and the original volume.

```
vxassist snapclear e: plex=volume1-01
```

This clears the association of e: (the snapshot volume) with its master volume.

```
vxassist snapclear F: snapvolume01
```

This command clears the association of master volume (F:) with its snapshot volume named "snapvolume01."

Note: When you use the `snapback` command to clear the association of the master volume, you must use the `<SnapPlexName>` or the `<Snapshot name>` when referring to the snapshot volume.

vxassist snapabort

```
vxassist [-o keepplex] [-g<DynamicDiskGroupName>] snapabort  
<VolumeName | DriveLetter> |VmName | DrivePath> [plex=<SnapPlexName>]
```

Aborts a snapshot mirror plex that has been attached to a volume through a prepare the volume for snapshot operation or a snapback operation by either deleting it or converting it to an ordinary mirror. In cases where the deleted snap plex is the last snap plex and the resulting volume is simple or striped, the `vxassist snapabort` command deletes the DCO log volume also.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxassist snapabort` command.

The following attributes apply:

<code>[-o keepplex]</code>	Converts the snap plex to a mirror plex instead of deleting it.
<code>-g<DynamicDiskGroupName></code>	Needed only if you use the internal volume name.
<code><VolumeName></code>	The path name of the original volume, such as <code>\Device\HarddiskDmVolumes\DynDskGrp1\volume1</code> .
<code><DriveLetter></code>	The drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

`plex=<SnapPlexName>` Name of a mirror plex that was created when a prepare the volume for snapshot process began. You would need this only if you already prepared the volume for snapshot previously and you prepare the volume for snapshot again, and thus you have multiple snap plexes and want to abort one of them.

A `<SnapPlexName>` can be the plex name (such as `Volume2-01`) or the GUID of the mirror plex. A GUID is a unique internal number assigned to the plex. To determine the GUID for a given plex, use the command `vxvol -v volinfo` for the mirrored volume that contains the plex.

Examples

```
vxassist snapabort E:
```

This command aborts the prepare the volume for snapshot operation on volume E.

```
vxassist -o keepplex snapabort H: plex=Concatvol11-03
```

This command aborts the prepare the volume for snapshot operation on volume H, but it converts the plex `Concatvol11-03` to a volume of the same type as the original volume. The reason the `SnapPlexName` had to be specified is that there were two snap plexes, because the `vxassist prepare` command had been invoked twice on volume H.

```
vxassist snapabort E: plex=volume1-0
```

This command aborts the prepare the volume for snapshot operation on volume E with the designated plex.

vxassist rescan

```
vxassist [-b] rescan
```

Rescans all the storage system objects on the computer, including disks, volumes, and file systems. The `-b` option runs the command in the background.

vxassist refresh

```
vxassist [-b] refresh
```

Refreshes the volume, drive letter, and system information. The `-b` option runs the command in the background.

vxassist resetbus

```
vxassist resetbus
```

Performs a SCSI bus reset on all SCSI buses.

vxassist version

```
vxassist version
```

Displays the installed version of the product and the build number.

vxassist (Windows-specific)

Assigns or changes a drive letter or drive path for a volume. In the command line interface, if you want a drive letter, you must always assign it.

vxassist assign

```
vxassist [-f] [-g<DynamicDiskGroupName>] assign <VolumeName |  
DriveLetter  
|VmName|DrivePath>DriveLetter={A|B|...|Z}|DrivePath=<DrivePath>
```

<code>-f</code>	Forces the command. Use this option with caution.
<code>-g<DynamicDiskGroupName></code>	Required only if the internal volume name or the internal disk name is used.
<code><VolumeName></code>	The path name of the original volume, such as <code>\Device\HarddiskDmVolumes\DynDskGrp1\Volume1</code> .
<code><DriveLetter></code>	The current drive letter of the volume.
<code><VmName></code>	Internal name of the volume; requires the use of the <code>-g</code> option (for example, <code>-gDG1 Volume1</code>).
<code><DrivePath></code>	Existing drive path for the volume.
<code>DriveLetter={A B ... Z}</code>	The drive letter to be assigned.
<code>DrivePath=<DrivePath></code>	The drive path to be assigned.

For more information on assigning a drive letter or drive path:

See [“Add, change, or remove a drive letter or path”](#) on page 117.

See [“Mount a volume at an empty folder \(Drive path\)”](#) on page 119.

If the volume already has a drive letter, assigning a drive letter replaces the existing drive letter. In addition, if you assign a drive path to a volume that already has an existing drive letter, the existing drive letter remains.

Example

```
vxassist assign \Device\HarddiskDmVolumes\DynDskGrp2\volume1
DriveLetter=E:
```

This command assigns drive letter E to
 \Device\HarddiskDmVolumes\DynDskGrp2\volume1.

vxassist unassign

```
vxassist [-f] [-g<DynamicDiskGroupName>] unassign <VolumeName |
DriveLetter |VmName> [DrivePath=<DrivePath>]
```

Removes a drive letter or drive path for a volume. The following attributes apply:

-f	Forces the command. Use this option with caution.
-g<DynamicDiskGroupName>	Required only if the internal volume name is used.
<VolumeName>	The path name of the original volume, such as \Device\HarddiskDmVolumes\DynDskGrp1\Volume1.
<DriveLetter>	The current drive letter of the volume.
<VmName>	Internal name of the volume; requires the use of the -g option (for example, -gDG1Volume1).
DrivePath=<DrivePath>	The drive path to be unassigned.

More information on unassigning a drive letter or drive path is available at the following:

See [“Add, change, or remove a drive letter or path”](#) on page 117.

See [“Mount a volume at an empty folder \(Drive path\)”](#) on page 119.

If a drive path and a drive letter are not specified, the command removes the drive letter. If you want to remove the drive path, you must specify it.

Example

```
vxassist unassign E:
```

This command removes drive letter E from the volume.

vxassist create_part

```
vxassist create_part <DiskName |p#c#t#l#> <Length>  
[DriveLetter={A|B|...|Z}|DrivePath=<DrivePath>]
```

Creates a primary partition of <Length> megabytes on <DiskName> and optionally assigns either a drive letter or drive path to it. To specify the partition size in KB, add K to the number (500K); similarly, add G for GB and T for TB.

See [“Create primary and extended partitions”](#) on page 135.

You can also replace the disk name by the p#c#t#l# parameter, where the #s correspond to port, channel, target, and LUN of the disk. Note that there is no internal disk name for a basic disk.

Example

```
vxassist create_part Harddisk0 100 DriveLetter=E
```

Creates a primary partition of 100 MB with drive letter E on Harddisk0.

vxassist create_ext

```
vxassist create_ext <DiskName |p#c#t#l#> <Length>
```

Creates an extended partition of <Length> megabytes on <DiskName>. To specify the extended partition size in KB, add K to the number (500K); similarly, add G for GB and T for TB. You can replace the disk name by the p#c#t#l# parameter, where the #s correspond to port, channel, target, and LUN of the disk.

Example

```
vxassist create_ext Harddisk0 100
```

Creates an extended partition of 100 MB on Harddisk0.

vxassist delete_ext

```
vxassist delete_ext <DiskName |p#c#t#l#>
```

Deletes the extended partition on <DiskName>. You can replace the disk name by the p#c#t#l# parameter, where the #s correspond to port, channel, target, and LUN of the disk.

Example

```
vxassist delete_ext Harddisk0
```


Deletes the extended partition on Harddisk0.

vxassist create_drive

```
vxassist create_drive <DiskName |p#c#t#l#> <Length>  
[DriveLetter={A|B|...|Z}]
```

Creates a logical drive of *<Length>* megabytes in the extended partition on *<DiskName>*. To specify the logical drive size in KB, add K to the number (500K); similarly, add G for GB and T for TB. The extended partition must have been created first.

See [“New logical drives”](#) on page 137.

You can also replace the disk name by the `p#c#t#l#` parameter, where the `#s` correspond to port, channel, target, and LUN of the disk.

Example

```
vxassist create_drive Harddisk0 100 DriveLetter=E
```

Creates a logical drive of 100 MB on Harddisk0. The logical drive has the drive letter E.

There is no separate `delete_part` command. The `vxassist delete` command, which removes a volume, can be used to delete primary partitions or logical drives. The `delete_ext` command is needed since an extended partition does not have a name to be used in the `vxassist delete` command.

vxassist taghost

Run the command `vxassist taghost site=<sitename>` to add a host to a site

```
vxassist taghost site=sitel
```

This command tags localhost to site `site1`.

See [“Administering hosts for site-aware allocation through CLI”](#) on page 453.

vxassist showtag

Run the `vxassist showtag` command to display the site name to which a host belongs:

```
vxassist showtag
```

The host is tagged to site `site1`

See [“Administering hosts for site-aware allocation through CLI”](#) on page 453.

vxassist untagghost

Use the `vxassist untagghost` command to remove a host from the site.

For example, run the command `vxassist untagghost` on a host to display:

```
vxassist untagghost
```

Successfully untagged host from its site.

See [“Administering hosts for site-aware allocation through CLI”](#) on page 453.

vxassist set sitetype

Use `vxassist set sitetype` command for setting the Sitetype volume property of a Siteless volume to either Site Confined or Site Separated.

```
vxassist -g<DynamicDiskGroupName> set  
<VolumeName|DriveLetter|VmName|DrivePath> sitetype={confined |  
separated}
```

See [“Administering volumes for site-based allocation ”](#) on page 456.

vxassist resetsitetype

Use the `vxassist resetsitetype` to convert a Site Confined or Site Separated volume to Siteless volume.

```
vxassist -g<DynamicDiskGroupName> resetsitetype  
<VolumeName|DriveLetter|VmName|DrivePath>
```

See [“Administering volumes for site-based allocation ”](#) on page 456.

vxevac

The `vxevac` utility moves subdisks from the disk that is specified by `<DiskName>` to one specified destination disk `<NewDiskName>`. If no `<NewDiskName>` argument is specified, any nonvolatile, nonreserved disk in the same dynamic disk group can be used as a destination disk. The destination disk must be in the same dynamic disk group.

```
vxevac [-o notrackalign>] [-g<DynamicDiskGroupName>] <DiskName  
|p#c#t#l#> [<NewDiskName>]
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxevac` command.

<code>-o notrackalign</code>	Disable track alignment when moving to destination disk.
<code>-g<DynamicDiskGroupName></code>	Required only if the internal disk name is used.
<code><DiskName></code>	The name of the disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk; the <code>#s</code> correspond to port, channel, target, and LUN of the disk.
<code>NewDiskName</code>	Name of the disk to which the subdisks are moved.

Example

```
vxevac Harddisk0 Harddisk3
```

Moves all the subdisks from Harddisk0 to Harddisk3.

vxsd

Command for moving, splitting, and joining a subdisk.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxsd` command.

<code>mv</code>	Moves a subdisk.
<code>split</code>	Splits a subdisk.
<code>join</code>	Joins a subdisk.

Typing the following sequence for each keyword brings up a description of its syntax.

```
vxsd <Keyword> -?
```

Use `vxvol volinfo` to display the subdisk names. Use `vxvol -v volinfo` to display the size and offset of the subdisks. The `-g<DynamicDiskGroupName>` option needs to be specified since `<SubDiskName>` is not distinct among different dynamic disk groups.

vxsd mv

```
vxsd [-o notrackalign] -g<DynamicDiskGroupName> [-f] [-b] mv  
<SubDiskName> [<DiskName | p#c#t#l#>]
```

This command moves the named subdisk to the disk that is specified by *<DiskName | p#c#t#l#>*. If a disk is not specified, the move operation tries to find free space within the dynamic disk group. The Move Subdisk function can be used for load balancing by moving a subdisk from a heavily accessed disk to a more lightly loaded one.

This command can also be used to move the part of a volume that is on a failed or missing disk to a healthy one. Moving a subdisk for redundant volumes (mirrored or RAID-5) uses the redundant data to recreate the subdisk on the healthy disk. However, for nonredundant volumes (concatenated or striped), the data cannot be recreated. If you perform a subdisk move, you lose data. You can avoid this if you repair the disk. Because of this possibility of data loss, Storage Foundation normally does not let you move the subdisk for nonredundant volumes.

You can disable track alignment for the operation by using the `-o notrackalign` option. You can also force the operation by using the `-f` option if you don't need the data any more or you can no longer recover the data. By doing so, you retains the volume structure, but there is no guarantee that the data will be recoverable. The `-b` option performs the command in the background.

Moving a subdisk from one disk to another involves a certain amount of time. The SmartMove feature can help reduce the time for this operation.

See [“SmartMove”](#) on page 93.

Example

```
vxsd -gDynDskGrp1 mv Disk1-01 Harddisk2
```

This command moves subdisk Disk1-01 to Harddisk2 (belonging to the DynDskGrp1 dynamic disk group).

```
vxsd -gDynDskGrp1 -f mv Disk1-01
```

This command forces the move of subdisk Disk1-01 (which is on a failed disk) to another disk with enough free space.

vxsd split

```
vxsd -g<DynamicDiskGroupName> -s<Size> split <SubDiskName>
```

Splits a subdisk into two separate subdisks.

The `-s` option is required to specify the `<Size>` in MB of the first of the two subdisks to be created. The second subdisk occupies the remaining space that is used by the original subdisk.

Example

```
vxsd -gDynDskGrp1 -s100 split Disk3-03
```

This command splits the subdisk Disk3-03 into two subdisks, the first of which is 100 MB.

vxsd join

```
vxsd -g<DynamicDiskGroupName> join <SubDiskName>...
```

Combines two or more existing subdisks into one subdisk. To join subdisks, the subdisks must be contiguous on the same disk.

Example

```
vxsd -gDynDskGrp1 join Disk3-03 Disk3-04
```

In this example, the command joins the two subdisks Disk3-03 and Disk3-04 into a new subdisk called Disk3-05.

vxstat

The `vxstat` utility has keywords and operands for displaying statistics for dynamic disks, dynamic volumes, and subdisks of dynamic volumes. It can also display statistics at specified intervals and for specified numbers of times.

Typing the following sequence brings up a description of its syntax:

```
vxstat <Keyword> -?
```

```
vxstat [-g<DynamicDiskGroupName>] [-s] [-d] [-v] [-i<Interval>]  
[-c<Count>] [<Object>...]
```

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group from which the records are selected. For example, <code>-gDynDskGrp1</code> .
---	--

[<Object>...]	An object can be a disk or a volume (specified by the volume name or drive letter). If no object is specified, the statistics apply to all the volumes on the managed server. Multiple objects are separated by a space.
[-s]	Displays the statistics for subdisks on the objects specified. If the object specified is a disk, all subdisks on that disk are selected. If the object specified is a volume, all subdisks on that volume are selected.
[-d]	Displays the statistics for disks on the objects specified.
[-v]	Displays the statistics for volumes on the objects specified.
[-i<Interval>]	Displays the statistics every <Interval> seconds. The default and minimum interval is 5 seconds between the instances that the statistics are displayed.
[-c<Count>]	Stops displaying the interval statistics after <Count> times.

If the `-i<Interval>` option is specified without the `-c<Count>` option, the statistics runs continuously with the interval specified. If neither the `-i<Interval>` option nor the `-c<Count>` option is specified, the statistics display once.

Summary statistics for each object are displayed in one-line output records, preceded by two header lines. The output line consists of blank-separated fields for the object type and object name.

Statistics Types

The statistics types are as follows:

- Read Requests/Second
The number of read requests per second for selected storage objects.
- Write Requests/Second
The number of write requests per second for selected storage objects.
- Read Blocks/Second
The amount of read request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the read requests made.
- Write Blocks/Second

The amount of write request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the write requests made.

- **Average Time/Read Block**
The average time in microseconds that it takes to process a read block from the time a request is made until the data is returned.
- **Average Time/Write Block**
The average time in microseconds that it takes to process a write block from the time a request is made until the data is returned.
- **Queue Depth**
The current number of read and write requests in the queue for selected disks. It does not apply to volumes and subdisks.

You can use these statistics to improve performance on your system.

See [“Statistics overview”](#) on page 229.

Examples

```
vxstat -s
```

Displays the statistics for all subdisks on a managed server.

```
vxstat -d Harddisk5 Harddisk6
```

Displays the statistics for Harddisk5 and Harddisk6

```
vxstat -s E:
```

Displays the statistics for the subdisks of volume E.

```
vxstat -i 10 -c 5 -d
```

Displays 5 sets of statistics at 10-second intervals for all the disks on the managed server.

```
vxstat -d F:
```

Displays the statistics for the disks that are associated with volume F.

```
vxstat -v Harddisk5
```

Displays the statistics for the volumes that are associated with Harddisk5.

```
vxstat -gDynDskGrp1 -v
```

Displays the statistics for all volumes in DynDskGrp1.

```
vxstat -s F: G: H: Harddisk7
```

Displays the subdisk statistics for volumes F, G, H, and for Harddisk7.

vxtask

The `vxtask` utility has keywords and operands to:

- Provide percentage completion information of a task running on a managed server.
Disk group and volume options let you specify a disk group or volume.
- Pause or resume a task
You can pause or resume a task only on a node where the task is running. In case of a CVM cluster, if a command is executed from a slave and a task begins on a master node, then the task can be paused only on the master node. The details of tasks that are paused or resumed are shown in the Alerts.

Note: You cannot pause or resume the following tasks:

- Volume shrink and volume format
 - RAID operations
 - DRL Resync
 - Resync operations when volume is in read-write back mode
 - Reclaim operations
 - Refresh and rescan operations
-

The `vxtask list /?` command brings up a description of the following syntax:

```
vxtask [-x] [-p] [-g<DynamicDiskGroupName>] [-v<VolumeName>]
|Driveletter | <VmName>] list
```

The following attributes can apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group for which the task information is displayed. For example, <code>-gDynDskGrp1</code> .
---	--

- [-v<VolumeName | Driveletter | VmName>]
Displays task information for the specified volume. The <VolumeName> is the path name of the volume. The <VmName> parameter requires the use of the -g option, such as -gDG1 Volume1.
- x
Displays the extended information such as, Task Name, Source, Target, TaskComplete(%), TaskID, and the State.

Each task has a unique task ID. The State displays if the task is Paused or is Running.
- p
Displays the tasks that are paused and the ones that can be paused, along with the extended information.

The `vxtask pause /?` command brings up a description of the following syntax:

```
vxtask [-g<DynamicDiskGroupName>]
[-v<VolumeName>|<Driveletter>|<VmName>] -t<TaskID> [-s] pause
```

Where,

- g<DynamicDiskGroupName>
Name of the dynamic disk group for which the task information is displayed. For example, -gDynDskGrp1.
- [-v<VolumeName>|<Driveletter>|<VmName>]
Displays task information for the specified volume. The <VolumeName> is the path name of the volume. The <VmName> parameter requires the use of the -g option, such as -gDG1 Volume1.
- t<TaskID>
Unique ID of a task that is used while pausing or resuming the task.
- s
Silent mode. Suppresses the user warnings and prompts.

The `vxtask resume /?` command brings up a description of the following syntax:

```
vxtask [-g<DynamicDiskGroupName>]  
[-v<VolumeName>|<Driveletter>|<VmName>] -t<TaskID> resume
```

Where,

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group for which the task information is displayed. For example, <code>-gDynDskGrp1</code> .
<code>[-v<VolumeName> <Driveletter> <VmName>]</code>	Displays task information for the specified volume. The <code><VolumeName></code> is the path name of the volume. The <code><VmName></code> parameter requires the use of the <code>-g</code> option, such as <code>-gDG1 Volume1</code> .
<code>-t<TaskID></code>	Unique ID of a task that is used while pausing or resuming the task.

By default the information on all the tasks running on the managed server are displayed unless the `-g` or `-v` options is specified.

Examples

```
vxtask -x list
```

Displays the task details and percentage completion information for all the tasks running on a managed server.

```
vxtask -x -gDynDskGrp1 list
```

Displays the details of the tasks relating to the disk group `DynDskGrp1` and its percentage completion information.

```
Vxtask -t TaskID pause
```

Pauses the specified task.

Note: If the task includes certain subtasks, you cannot pause them individually. The entire set of task (parent task) is paused.

```
Vxtask -t Taskid resume
```

Resumes the specified task.

You can resume only the tasks that are in Paused state. The operation resumes from the point where it was paused. However, if the system restarts or the resource fails over when the task is paused, the task is reinitialized.

vxedit

The `vxedit` command has keywords to associate a comment with a dynamic disk group, disk, volume, plex, or subdisk, or to set an alias for an enclosure.

<code>set comment</code>	Associates a comment with a dynamic disk group, disk, volume, plex, or subdisk.
<code>set alias</code>	Set alias for an enclosure.
<code>rename</code>	Changes the name of a volume, plex, or disk media.

Typing the following sequence brings up a description of its syntax:

```
vxedit <Keyword> -?
```

vxedit set comment

```
vxedit [-g<DynamicDiskGroupName>] set comment=<comment>[<Object>]
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxedit set comment` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group for which the task information is displayed. For example, <code>-gDynDskGrp1</code> .
<code><comment></code>	The comment text. The maximum size of comment is 40 bytes.
<code><Object></code>	Can be a dynamic volume, disk, plex, disk media, subdisk, or enclosure. Plex and subdisk require use of the <code>-g</code> option. If no object is specified, the command defaults to dynamic disk group (requires use of the <code>-g</code> option).

This command allows the user to add comments to any dynamic disk group, disk, volume, plex, or subdisk.

Example

```
vxedit set comment=BootDisk Harddisk0
```

Set the comment field of Harddisk0 to "BootDisk."

vxedit set alias

```
vxedit [-g<DynamicDiskGroupName>] set
alias=<EnclosureNewAlias><EnclosureName | EnclosureOldAlias>
```

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group for which the task information is displayed. For example, <code>-gDynDskGrp1</code> .
<code><EnclosureNewAlias></code>	New alias name for enclosure.
<code><EnclosureName EnclosureOldAlias></code>	Name of enclosure or old alias of enclosure.

This command allows the user to set an alias for an enclosure.

Example

```
vxedit set alias=ABC1 DEC0
```

Set the alias for enclosure DEC0 to ABC1.

vxedit rename

```
vxedit [-g<DynamicDiskGroupName>] rename
<OldVolumeName><NewVolumeName> [-f] | <OldPlexName><NewPlexName> |
<OldDMName><NewDMName>
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxedit rename` command.

The following attributes apply:

<code>-g<DynamicDiskGroupName></code>	Name of the dynamic disk group for which the task information is displayed. For example, <code>-gDynDskGrp1</code> .
<code><OldVolumeName></code>	Old name for volume.
<code><NewVolumeName></code>	New name for volume.
<code>-f</code>	Forces the rename of the volume.
<code><OldPlexName></code>	Old name for plex.
<code><NewPlexName></code>	New name for plex.

<code><OldDMName></code>	Old name for disk media.
<code><NewDMName></code>	New name for disk media.

This command allows the user to change the name of a volume, plex, or disk media to a new name.

Example

```
vxedit rename Volume1 VolumeABC
```

Rename Volume1 to VolumeABC.

vxunreloc

Reverses the process of hot relocation or evacuation and moves the hot-relocated subdisks back onto the disk that was replaced after a disk failure.

```
vxunreloc [-g<DynamicDiskGroupName>] [-f] <DiskName |p#c#t#l#>
```

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxunreloc` command.

Note: Dynamic disks belonging to a SAN Disk Group do not support the `vxunreloc` command.

<code>-g<DynamicDiskGroupName></code>	Required only if the internal disk name is used.
<code>-f</code>	Unrelocates a subdisk to a different offset if unrelocating to the original offset is not possible.
<code>DiskName</code>	Name of the disk that was relocated to another disk. For example, <i>device name</i> such as <code>Harddisk2</code> , or <i>internal disk name</i> such as <code>Disk2</code> (requires use of <code>-g</code> option, for example, <code>-gDG1 Disk2</code>).
<code>p#c#t#l#</code>	Another way to indicate a disk name; the <code>#s</code> correspond to port, channel, target, and LUN of the disk.

All the disks that are involved in the process must be in the same dynamic disk group.

Example

```
vxunreloc Harddisk0
```

This command unrelocates Harddisk0, which had been previously relocated to another hard disk.

See [“Hot relocation”](#) on page 216.

vxdmpadm

The `vxdmpadm` utility provides commands for Dynamic Multi-Pathing (DMP DSMs).

<code>dsminfo</code>	Displays the information about the DSM.
<code>arrayinfo</code>	Displays the information about an array.
<code>deviceinfo</code>	Displays the information about a device (that is, a hard disk in an array).
<code>pathinfo</code>	Provides the information about a path.
<code>arrayperf</code>	Displays the performance statistics about an array.
<code>deviceperf</code>	Displays the performance statistics about a device.
<code>pathperf</code>	Displays the performance statistics about the paths that are connected to a device.
<code>allperf</code>	Displays the performance statistics for all paths for all devices.
<code>iostat</code>	Displays the I/O statistics for a single disk, all disks in an array, or all the disks of all the arrays in the host.
<code>cleardeviceperf</code>	Resets the performance I/O statistics of a device.
<code>cleararrayperf</code>	Resets the performance I/O statistics of an array that contains the specified disk.
<code>clearallperf</code>	Resets the performance I/O statistics of all arrays under the control of the DMP DSMs that the specified disk belongs to.
<code>setdsmscsi3</code>	Enables or disables SCSI3 support for the DMP DSMs that control the specified disk. On Windows Server operating systems, can enable or disable SCSI3 support for Veritas DSMs that are installed but with storage not yet connected.
<code>setarrayscsi3</code>	Enables or disables SCSI3 support in the registry of the array that contains the specified disk.
<code>setattr dsm</code>	Sets DSM attributes.

setattr array	Sets the array attributes.
setattr device	Sets the device attributes.
setattr path	Sets the path attributes.
set isislog	Alters the trace level of the VEA/ISIS log dynamically.
rescan	Rescans all the storage system objects on the computer.
disk list	Lists the details of all the disks on the managed server.
getdsmattrib	Displays the information for installed Veritas DSMs with no storage connection.
getmpioparam	Displays MPIO parameters for installed Veritas DSMs
setmpioparam	Sets MPIO parameters for installed Veritas DSMs

A full explanation of the DMP DSMs software and commands is available.

See [“DMP overview”](#) on page 399.

Each keyword is followed by the volume name or drive letter. Typing the following sequence for each keyword brings up a description of its syntax:

```
vxdmpadm<Keyword> -?
```

vxdmpadm dsminfo

```
vxdmpadm [-g<DynamicDiskGroupName>] dsminfo <DiskName> | p#c#t#l#>
```

This command displays the DSM name, arrays in the DSM, and the array load balance policy. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG1 Disk2`). The `#s` in the `p#c#t#l#` parameter correspond to the port, channel, target, and LUN of a disk.

Example

```
vxdmpadm dsminfo Harddisk5
```

Displays the DSM information in which `Harddisk5` participates.

vxdmpadm arrayinfo

```
vxdmpadm [-g<DynamicDiskGroupName>] arrayinfo <DiskName> | p#c#t#l#>
```

This command displays the array name, the devices in the array, and the array load balance policy. It also displays the tunable parameters (control timer settings) that

affect the testing and failover of paths. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The internal disk name must be preceded by the `-g` option (for example, `-gDG1 Disk2`). The `#s` in the `p#c#t#l#` parameter correspond to the port, channel, target, and LUN of a disk.

Example

```
vxdmpadm arrayinfo Harddisk5
```

Displays the array information for the array in which `Harddisk5` participates.

vxdmpadm deviceinfo

```
vxdmpadm [-g<DynamicDiskGroupName>] deviceinfo <DiskName> |  
p#c#t#l#>...
```

This command displays the device name, the internal disk name, number of paths, status, and load balance policy. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG1 Disk2`). The `p#c#t#l#` parameter corresponds to the port, channel, target, and LUN of a disk.

Example

```
vxdmpadm deviceinfo Harddisk5 Harddisk6
```

Displays DMP DSMs related information about `Harddisk5` and `Harddisk6`.

vxdmpadm pathinfo

```
vxdmpadm [-g<DynamicDiskGroupName>] pathinfo <DiskName> | p#c#t#l#>...
```

This command displays path status, load balance policy, port, target, and LUN along with device name, internal disk name, and number of paths. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG1 Disk2`). The `p#c#t#l#` parameter corresponds to the port, channel, target, and LUN of a disk.

Example

```
vxdmpadm pathinfo Harddisk5 Harddisk6
```

Displays path information for `Hardisk5` and `Harddisk6`.

vxdmpadm arrayperf

```
vxdmpadm [-g<DynamicDiskGroupName>] arrayperf <DiskName> | p#c#t#l#>
```

This command displays the date and time of the performance statistics, the array name, all the devices in the array, the paths for each device, and the performance statistics of all the devices in the array. The statistics are grouped by each path that is connected to a device. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The internal disk name must be preceded by the `-g` option (for example, `-gDG1 Disk2`). The `#s` in the `p#c#t#l#` parameter corresponds to the port, channel, target, and LUN of a disk.

The performance statistics represent the aggregate workload of the array.

It includes the following number of:

- Reads
- Writes
- Bytes read
- Bytes written

Example

```
vxdmpadm arrayperf Harddisk5
```

Displays the performance statistics of the devices in the array in which `Harddisk5` participates.

vxdmpadm deviceperf

```
vxdmpadm [-g<DynamicDiskGroupName>] deviceperf <DiskName> |  
p#c#t#l#>...
```

This command displays the date and time of the performance statistics, the device name, the paths for the device, and the performance statistics of the device. The statistics are grouped by each path that is connected to the device. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG1 Disk2`). The `p#c#t#l#` parameter corresponds to the port, channel, target, and LUN of a disk.

The performance statistics represent the aggregate workload of the device.

It includes the following number of:

- Reads

- Writes
- Bytes read
- Bytes written

Example

```
vxdmpadm deviceperf Harddisk6
```

Displays the performance statistics of Harddisk6.

vxdmpadm pathperf

```
vxdmpadm [-g<DynamicDiskGroupName>] pathperf <DiskName> | p#c#t#l#>...
```

This command displays the date and time of the performance statistics, the device name, and the performance statistics of the device. The statistics are grouped by each path that is connected to the device. The *DiskName* parameter can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG1 Disk2`). The `p#c#t#l#` parameter corresponds to the port, channel, target, and LUN of a disk.

The performance statistics represent the workload of the paths that are connected to the device.

It includes the following number of:

- Reads
- Writes
- Bytes read
- Bytes written

Example

```
vxdmpadm pathperf Harddisk6
```

Displays the performance statistics of Harddisk6.

vxdmpadm allperf

```
vxdmpadm allperf
```

This command displays performance information of the paths for all devices.

For devices, this includes the following:

- Name of the counter

- Path name
- Device name
- Array name
- Counter-value

For path summaries, this includes the following:

- Name of the counter
- Path name
- Array name
- Counter-value

vxdmpadm iostat

This command displays the I/O statistics for a single disk, all disks in an array, or all the disks of all the arrays in the host. You can group the statistics by device or array. If no `groupby` is specified, then statistics is shown for all the paths. The statistics can be set to display after a specified number of seconds and for a specified number of times through this CLI option.

```
vxdmpadm iostat vxdmpadm [-g<DynamicDiskGroupName>] iostat
[showdevice|showarray] [interval=#] [count=#] <DiskName> | p#c#t#l#>
vxdmpadm [-g<DynamicDiskGroupName>] iostat [showdevice | showarray
| showall] [groupby=<device | array>] [-u <KB | MB | GB>]
[interval=seconds] [count=N] [<diskname | p#c#t#l#>]
```

The following attributes apply:

`showdevice|showarray`

Display I/O statistics:

- `showdevice` specifies paths of the disk. (Default is `showdevice`.)
- `showarray` specifies all the paths of the array that contains the disk.

`showall`

Displays the I/O statistics for all the paths of all the arrays that are connected to the host. `Diskname` should not be specified with this option.

`groupby=device`

The I/O statistics is shown per device.

If no `groupby` is specified, then statistics is shown for all devices.

groupby=array	The I/O statistics is shown per array
KB	Data that is read or written is shown in KiloBytes/Second
MB	Data that is read or written is shown in MegaBytes/Second
GB	Data that is read or written is shown in GigaBytes/Second
interval=#	# specifies the amount of time (seconds) between displays of the I/O statistics. (Default value is 5 seconds between displays.)
count=#	# specifies the number of times the I/O statistics are displayed. (Default is infinite number of displays.)
<DiskName>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
p#c#t#l#	The #s correspond to the port, channel, target, and LUN of a disk.

vxdmpadm cleardeviceperf

```
vxdmpadm [-g<DynamicDiskGroupName>] cleardeviceperf
```

```
<DiskName> | p#c#t#l#>
```

This command resets the performance I/O statistics of a device.

The following attributes apply:

<DiskName>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
p#c#t#l#	The #s correspond to the port, channel, target, and LUN of a disk.

vxdmpadm cleararrayperf

```
vxdmpadm [-g<DynamicDiskGroupName>] cleararrayperf<DiskName> |  
p#c#t#l#>
```

This command resets the performance I/O statistics of an array that contains the specified disk.

The following attributes apply:

<DiskName> The designated disk, which can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG2 Disk1`).

p#c#t#l# The #s correspond to the port, channel, target, and LUN of a disk.

vxdmpadm clearallperf

```
vxdmpadm [-g<DynamicDiskGroupName>] clearallperf<DiskName> | p#c#t#l#>
```

This command resets the performance I/O statistics of all arrays under the control of the DMP DSMs that the specified disk belongs to.

The following attributes apply:

<DiskName> The designated disk, which can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG2 Disk1`).

p#c#t#l# The #s correspond to the port, channel, target, and LUN of a disk.

vxdmpadm setdsmscsi3

```
vxdmpadm [-g<DynamicDiskGroupName>] setdsmscsi3 scsi3support=0|1  
<DiskName> | p#c#t#l#> | [dsmName=ALL | <DsmFileName>]
```

This command enables or disables SCSI3 support in the registry of the DMP DSMs that control the specified disk.

The following attributes apply:

scsi3support=0|1

- 0 disables SCSI3 support in the DMP DSMs
- 1 enables SCSI3 support in the DMP DSMs

<i><DiskName></i>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
<code>p#t#l#</code>	The <code>#s</code> correspond to the port, channel, target, and LUN of a disk.
<code>dsmName=ALL <DsmFileName></code>	This attribute applies the <code>scsi3support=0 1</code> selection to enable or disable SCSI3 on installed DSMs before a connection is configured. You can apply to all installed DSMs or to a particular DSM, specified by file name. You can use the <code>vxdmpadm getdsmattrib</code> command to get a list of DSM file names.

vxdmpadm setarrayscsi3

```
vxdmpadm [-g<DynamicDiskGroupName>] setarrayscsi3 scsi3support=0|1
<DiskName> | p#c#t#l#>
```

This command enables or disables SCSI3 support in the registry of the array that contains the specified disk.

The following attributes apply:

<code>scsi3support=0 1</code>	<ul style="list-style-type: none"> ■ 0 disables SCSI3 support in the DMP DSMs ■ 1 enables SCSI3 support in the DMP DSMs
-------------------------------	---

<i><DiskName></i>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
-------------------------	--

<code>p#c#t#l#</code>	The <code>#s</code> correspond to the port, channel, target, and LUN of a disk.
-----------------------	---

vxdmpadm setattr dsm

```
vxdmpadm [-g<DynamicDiskGroupName>] setattr dsm
[loadbalancepolicy=FO|RR|LQ|LB|BP [blockshift=#]] [primarypath=#]
[testpathretrycount=# scsicmdtimeout=# kernalsleeptime=#
failoverretrycount=#] <DiskName> | p#c#t#l#> [dsmName=ALL |
<DsmFileName>]
```

This command sets the load balance policy and primary path of the DSM to which the designated disk belongs. It also lets you set tunable parameters (control timer

settings) that affect the testing and failover of the paths. The following attributes apply:

<code>loadbalancepolicy=</code>	Specifies the load balance policy for the DSM where: <ul style="list-style-type: none">■ FO specifies Fail Over Only (Active/Passive)■ RR specifies Round Robin (Active/Active)■ LQ specifies Dynamic Least Queue Depth■ LB specifies Least Blocks■ BP specifies Balanced Path
<code>FO RR LQ </code> <code>LB BP</code>	
<code>blockshift=#</code>	# specifies the number of contiguous I/O blocks that are sent along a path to an Active/Active array before switching to the next available path. (Default value is 2048 blocks.)
<code>primarypath=#</code>	# specifies the primary path of the DSM. For example, <code>primarypath=1-1-1</code> sets path 1-1-1 as the primary path of the DSM.
<code>testpathretrycount=#</code>	Range: 0 ~ 50 times. Default: 0.
<code>scsicmdtimeout=#</code>	# specifies the amount of time a SCSI command waits for a device to respond to a SCSI command. Default is 30 seconds. Maximum is 120 seconds.
<code>kernalsleeptime=#</code>	# specifies the interval of time the DMP DSMs kernel waits between attempts to recover a path from an error condition. Default is 200 milliseconds. Maximum is 10000 milliseconds.
<code>failoverretrycount=#</code>	# specifies the number of times DMP DSMs attempts a failover to another path when the current path fails. Default is 0. Maximum is 50 times. Note: For DMP DSMs, Failover Retry Count does not apply to the EMC Symmetrix array.
<code><DiskName></code>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
<code>p#c#t#l#</code>	The #s correspond to the port, channel, target, and LUN of a disk.

dsmName = ALL | <DsmFileName> Specifies whether to apply settings to all installed DSMs or only to the DSM specified by file name. Only the load balance policy settings are applied. If other settings are specified, the command returns an error.

Note: This parameter can be applied with or without a storage connection. If the array hasn't been configured with load balance settings previously, it inherits the load balance settings that are specified for the DSM once it is connected. However, if the array was previously connected and has been configured with load balance settings in the registry, the previous settings are maintained when the array is reconnected.

Example

```
vxdmpadm setattr dsm loadbalancepolicy=FO
primarypath=1-1-0 scsicmdtimeout=34 Harddisk6
```

Sets the load balance policy of the DSM to Fail Over Only (Active/Passive), the primary path to path 1-1-0, the scsicmdtimeout to 34 seconds.

vxdmpadm setattr array

```
vxdmpadm [-g<DynamicDiskGroupName>] setattr array
[loadbalancepolicy=FO|RR|RS [path#=state#] |LQ|WP
[path#=weight#]|LB|BP [blockshift=#]] [primarypath=#]
[testpathretrycount=# scsicmdtimeout=# kernalsleeptime=#
failoverretrycount=# ]<DiskName> | p#c#t#l#>
```

This command sets the load balance policy and primary path of the array to which the designated disk belongs. It also lets you set tunable parameters (control timer settings) that affect the testing and failover of the paths. The following attributes apply:

<code>loadbalancepolicy=</code>	Specifies the load balance policy for the array where: <ul style="list-style-type: none">■ FO specifies Fail Over Only (Active/Passive)■ RR specifies Round Robin (Active/Active)■ RS specifies Round Robin with Subset■ LQ specifies Dynamic Least Queue Depth■ WP specifies Weighted Paths■ LB specifies Least Blocks■ BP specifies Balanced Path <p>Note: Round Robin with Subset and Weighted Paths are available for A/PC and ALUA arrays only at the device level. They are not available at the array level.</p>
<code>FO RR RS</code>	
<code> LQ WP</code>	
<code> LB BP</code>	
<code>path#=state#</code>	<code>state#</code> specifies either standby (0) or active (1) state of the specified <code>path#</code> . For example, 1-1-1=0 means that path 1-1-1 is assigned a state of standby. (Default state of a path is active (1).)
<code>path#=weight#</code>	<code>weight#</code> specifies the weight that is assigned to the specified <code>path#</code> . For example, 1-1-1=10 means that path 1-1-1 is assigned a weight of 10. The path with the lowest weight indicates the most-favored path for I/O. (Range of values for <code>weight#</code> : 0 - 255.)
<code>blockshift=#</code>	<code>#</code> specifies the number of contiguous I/O blocks that are sent along a path to an Active/Active array before switching to the next available path. (Default value is 2048 blocks.)
<code>primarypath=#</code>	<code>#</code> specifies the primary path of the array. For example, <code>primarypath=1-1-1</code> sets path 1-1-1 as the primary path of the array.
<code>testpathretrycount=#</code>	<code>#</code> specifies the number of times the testpath routine retries to test the health of the path. Range: 0 ~ 50 times. Default: 0.
<code>scsicmdtimeout=#</code>	<code>#</code> specifies the amount of time a SCSI command waits for a device to respond to a SCSI command. Range: 30 ~ 120 seconds. Default is 30 seconds.
<code>kernalsleeptime=#</code>	<code>#</code> specifies the interval of time the DMP DSMs kernel waits between attempts to recover a path from an error condition. Range: 200 ~ 10000 milliseconds. Default is 200 milliseconds.

`failoverretrycount=#` # specifies the number of times the failover routine retries to failover the path. Range: 0 ~ 50 times. Default: 0.

Note: For DMP DSMs, Failover Retry Count does not apply to the EMC Symmetrix array.

`<DiskName>` The designated disk, which can be specified by the device name (such as `Harddisk2`) or the internal disk name (such as `Disk2`). The `-g` option must precede the internal disk name (for example, `-gDG2 Disk1`).

`p#c#t#l#` The #s correspond to the port, channel, target, and LUN of a disk.

Examples

```
vxdmpadm setattr array loadbalancepolicy=FO primarypath=2-1-0
testpathretrycount=4 Harddisk6
```

Sets the load balance policy of the array to Fail Over Only (Active/Passive), the primary path to path 2-1-0, the testpathretrycount to 4 times.

```
vxdmpadm setattr array loadbalancepolicy=RR Harddisk6
```

Sets the load balance policy of the array to Round Robin (Active/Active).

vxdmpadm setattr device

```
vxdmpadm [-g<DynamicDiskGroupName>] setattr device
[loadbalancepolicy=FO|RR|RS [path#=state#] |LQ|WP
[path#=weight#]|LB|BP [blockshift=#]] [primarypath=#] <DiskName> |
p#c#t#l#>
```

This command sets the load balance policy and primary path of the device to which the designated disk belongs. The following attributes apply:

`loadbalancepolicy=` Specifies the load balance policy for the device where:

- | | |
|------------------------|--|
| <code>FO RR RS </code> | <ul style="list-style-type: none"> ■ FO specifies Fail Over Only (Active/Passive) ■ RR specifies Round Robin (Active/Active) ■ RS specifies Round Robin with Subset |
| <code>LQ WP LB</code> | <ul style="list-style-type: none"> ■ LQ specifies Dynamic Least Queue Depth ■ WP specifies Weighted Paths ■ LB specifies Least Blocks ■ BP specifies Balanced Path |

<code>path#=state#</code>	<code>state#</code> specifies either standby (0) or active (1) state of the specified <code>path#</code> . For example, 1-1-1=0 means that path 1-1-1 is assigned a state of standby. (Default state of a path is active (1).)
<code>path#=weight#</code>	<code>weight#</code> specifies the weight that is assigned to the specified <code>path#</code> . For example, 1-1-1=10 means that path 1-1-1 is assigned a weight of 10. The path with the lowest weight indicates the most-favored path for I/O. (Range of values for <code>weight#</code> : 0 - 255.)
<code>blockshift=#</code>	<code>#</code> specifies the number of contiguous I/O blocks that are sent along a path to an Active/Active array before switching to the next available path. (Default value is 2048 blocks.)
<code>primarypath=#</code>	<code>#</code> specifies the primary path of the device. For example, <code>primarypath=1-1-1</code> sets path 1-1-1 as the primary path of the device.
<code><DiskName></code>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
<code>p#c#t#l#</code>	The <code>#s</code> correspond to the port, channel, target, and LUN of a disk.

Example

```
vxddmpadm setattr device loadbalancepolicy=FO primarypath=1-1-0
Harddisk6
```

Sets the load balance policy of the device to Fail Over Only (Active/Passive) and the primary path to path 1-1-0 for Harddisk6.

vxddmpadm setattr path

```
vxddmpadm [-g<DynamicDiskGroupName>] setattr path
[loadbalancepolicy=FO|RR|RS [path#=state#] |LQ|WP
[path#=weight#]|LB|BP [blockshift=#]] [primarypath=#] <DiskName> |
p#c#t#l#>...
```

This command sets the load balance policy and primary path of the designated disk. The following attributes apply:

<code>loadbalancepolicy=</code>	Specifies the load balance policy for the device where: <ul style="list-style-type: none">■ FO specifies Fail Over Only (Active/Passive)■ RR specifies Round Robin (Active/Active)■ RS specifies Round Robin with Subset■ LQ specifies Dynamic Least Queue Depth■ WP specifies Weighted Paths■ LB specifies Least Blocks■ BP specifies Balanced Path
<code>FO RR RS </code> <code>LQ WP LB</code>	
<code>path#=state#</code>	<code>state#</code> specifies either standby (0) or active (1) state of the specified <code>path#</code> . For example, 1-1-1=0 means that path 1-1-1 is assigned a state of standby. (Default state of a path is active (1).)
<code>path#=weight#</code>	<code>weight#</code> specifies the weight that is assigned to the specified <code>path#</code> . For example, 1-1-1=10 means that path 1-1-1 is assigned a weight of 10. The path with the lowest weight indicates the most-favored path for I/O. (Range of values for <code>weight#</code> : 0 - 255.)
<code>blockshift=#</code>	<code>#</code> specifies the number of contiguous I/O blocks that are sent along a path to an Active/Active array before switching to the next available path. (Default value is 2048 blocks.)
<code>primarypath=#</code>	<code>#</code> specifies the primary path of the device. For example, <code>primarypath=1-1-1</code> sets path 1-1-1 as the primary path of the device.
<code><DiskName></code>	The designated disk, which can be specified by the device name (such as <code>Harddisk2</code>) or the internal disk name (such as <code>Disk2</code>). The <code>-g</code> option must precede the internal disk name (for example, <code>-gDG2 Disk1</code>).
<code>p#c#t#l#</code>	The <code>#s</code> correspond to the port, channel, target, and LUN of a disk.

Example

```
vxdmpadm setattr path loadbalancepolicy=FOprimarypath=1-1-0  
Harddisk6
```

Sets the load balance policy of the device to Fail Over Only (Active/Passive) and the primary path to path 1-1-0 for Harddisk6.

vxdmpadm set isislog

```
vxdmpadm set isislog level=<#|default> [reg=<y|n>]
```

This command alters trace level of the VEA/ISIS log dynamically. The following attributes apply:

- `level=#` Specifies the trace level to be set as current trace level, where #can be:
- 1 specifies the trace level Info1
 - 2 specifies the trace level Info2
 - 3 specifies the trace level Info3
 - 4 specifies the trace level Info
 - 5 specifies the trace level Warning
 - 6 specifies the trace level Error (default)
- `reg=y|n` Specifies whether to update the trace level in the registry or not.
- Specify `reg=y` to update the trace level in the registry.
 - Specify `reg=n` to not update the trace level in the registry.
- By default, the trace level is not updated in the registry if this attribute is not specified.

Example

```
vxdmpadm set isislog level=default
```

Sets the trace level to 'Error'(6) and does not update the trace level in the registry.

```
vxdmpadm set isislog level=4 reg=y
```

Sets the trace level to 'Info'(4) and updates the trace level in the registry.

vxdmpadm rescan

```
vxdmpadm [-b] rescan
```

Rescans all the storage system objects on the computer, including disks, volumes, and file systems. The `-b` option runs the command in the background.

vxdmpadm disk list

```
vxdmpadm disk [-g<DynamicDiskGroupName>] [-v] list
```

Lists device name, internal disk name, disk group, disk partition style, size, free space, LUN, status, serial number, ThinProvisioningType (ThinNoReclaim or

ThinReclaim for thin provisioned disks), ProvisionedSize (current size of thin provisioned disks), etc. for all the disks on the managed server.

The `-g` option limits the list of disks to the specified disk group. The `-v` (verbose) option displays signature, status, bus type, port, target, channel, LUN, track alignment offset, and track alignment setting information for the disks in the list.

Note: Not all manufacturers of thin provisioned disks enable the display of the current provisioned size.

Example

```
vxdmpadm disk -gDB1 -v list
```

Gives information on all the disks in the DB1 disk group. The `-v` option includes the signature, status, bus type, port, channel, target, and LUN information.

```
vxdmpadm disk [-g<DynamicDiskGroupName>] [-v] list
```

vxdmpadm getdsmattrib

```
vxdmpadm getdsmattrib [dsmName=ALL | <DsmFileName>]
```

ALL is the default; it gets the information for all installed Veritas DSMs.

<DsmFileName>: gets the information only for the DSM with that file name.

This command displays information for Veritas DSMs that have been installed in the system, with no storage connection required, as follows:

- The Veritas DSM file name
- The Veritas DSM full name which has been registered with MPIO
- The current load balance of the Veritas DSM
- The SCSI3 support setting of the Veritas DSM

vxdmpadm getmpioparam

```
vxdmpadm getmpioparam [dsmName=ALL | <DsmFileName>]
```

If a DSM file name is specified, the command displays MPIO parameter settings for the specified DSM. Otherwise, the command displays MPIO parameter settings for all Veritas DSMs that have been installed on a system.

vxdmpadm setmpioparam

```
vxdmpadm setmpioparam [default] |[PathVerifyEnabled=0|1]
[PathVerificationPeriod=#][PDORemovePeriod=#] [RetryCount=#]
[RetryInterval=#] [dsmName=ALL | <DsmFileName>]
```

This command enables changing MPIO timer parameters or restoring default MPIO timer parameters.

Warning: Veritas DSMs always use the default MPIO timer parameters. Therefore, this command should be used only when Microsoft has advised to change the settings for debugging purposes. Attempting to change these MPIO parameters would affect the behavior of a DSM for I/O error, path failover, and DSM performance.

default	Restores the default settings of these MPIO parameters
PathVerifyEnabled=0 1	Default: 0 (not enabled) If enabled (value 1) MPIO performs path verification every <i>N</i> seconds where <i>N</i> is the value set in PathVerificationPeriod.
PathVerificationPeriod=#	# specifies the amount of time MPIO performs the path verification, if path verification is enabled. Default: 30 seconds. Maximum value: MAXULONG.
PDORemovePeriod=#	# specifies the amount of time an MPIO pseudo LUN, which represents a disk under MPIO control, stays in memory after all paths of a disk have been removed from the system. It also specifies how long the pending I/O should fail after all paths have been removed from a disk. Default: 20 seconds. Maximum value: MAXULONG.
RetryCount=#	# specifies the number of times DSM asks MPIO to retry the I/O when an I/O error occurs. Default: 3. There is no maximum value.
RetryInterval=#	# specifies the amount of time MPIO should retry a failed I/O. Default: 1 second.
dsmName = ALL <DsmFileName>	Specifies whether to apply these settings to all installed DSMs or only to the DSM specified by file name. Note: You can use the <code>vxdmpadm getdsmattrib</code> command if you want to find out the file names of installed DSMs.

vxcbr

The `vxcbr` utility provides the ability for users to back up and restore their Storage Foundation configuration. This utility does not back up and restores data. Only the backup of Storage Foundation configuration—that is, the disk group and logical volume layout on a server is taken.

The history of backup and restore operations is saved in the command line log, `vxcli.log`, at the following location:

```
C:\Program Files\Veritas\Veritas Volume Manager\logs
```

<code>backup</code>	Backs up the Storage Foundation configuration.
<code>restore</code>	Takes a latest backup and restores the Storage Foundation configuration.
<code>write_signature</code>	Manually writes a signature to a disk.

Typing the following sequence for each keyword brings up a description of its syntax:

```
vxcbr <Keyword> -?
```

Note: The `vxcbr` utility does not support Microsoft Disk Management Disk Groups. To back up the configuration of a Microsoft Disk Management Disk Group, use the Windows Server Backup feature of Windows.

vxcbr backup

```
vxcbr [-a] [-p<Path>] backup [<DynamicDiskGroupName>...]
```

Takes a backup of the SFW configuration and creates a text file providing the configuration details. The backup and the text file is saved in the product installation directory (normally `C:\ProgramData\Veritas\VXCBR`) or to a directory that is indicated in the `-p<Path>` parameter. The `vxcbr` utility creates a back up of only the configurations of disk groups that are online. Any offline or deported disk groups are skipped by this utility for backup purposes.

The following attributes apply:

<code>[-a]</code>	Indicates that all disk groups on the server are backed up. This option is recommended.
<code>-p<Path></code>	Specifies the path to the directory where the configuration information are archived. The default path is <code>%ALLUSERSPROFILE%\VERITAS\VXCBR</code> .

<DynamicDiskGroupName>... Limits the configuration to the disk group or disk groups indicated.

For example,

Run the following command to take a backup of a particular disk group:

```
vxcb backup DiskGroupName
```

Run the following command to take a backup of all the disk groups:

```
vxcb -a backup
```

This backs up configuration of all online SFW disk groups on the managed server.

vxcb restore

```
vxcb [-i] [-p<Path>] restore [<DynamicDiskGroupName>...]
```

Takes a latest backup and then restores the configuration. A folder *RestoreTimeStamp\DynamicDiskGroupName* is created at the specified backup path to save the latest backup.

The restore command only restores the disk group configurations to their original disks. The program assumes that you have not changed or replaced any of the ORIGINAL disks. However, if a disk is missing, you can specify the *-i* parameter to restore a configuration of an incomplete disk group. Probably, the only reason you might want to do an incomplete restore is to be able to access existing disks. An incomplete restore of the configuration is not recommended.

The following attributes apply:

<i>-i</i>	Indicates that you want to restore an incomplete configuration—that is, one of the disks in the configuration is missing. It is a forced restore of an incomplete disk group.
<i>-p<Path></i>	The path to where configuration information is saved.
<i><DynamicDiskGroupName>...</i>	Restores the disk group or groups specified.

Example

```
vxcb restore Dg01 Dg02
```

Restores the disk group configurations for Dg01 and Dg02.

Note: After you restore the configuration you must rescan the disks and then import the DGs.

vxabr write_signature

```
vxabr [-f<FileName>] write_signature [<<DiskName>=<DiskSignature>>...]
```

If the signature on a disk is no longer readable, you must manually add the signature back to the disk. This ensures that the configuration information is restored to the disk. You can find the signature information from a backup of the disk. If you have one or two disks, you can indicate the disk names and signatures in the command line. However, if you have multiple disks, you may find it easier to create a separate file with the disk names and the corresponding disk signatures. The utility reads the information from a specified file.

The following attributes apply:

[-f<FileName>]

Name of the file that has disk names and signatures for each disk you want the write signature to be applied to. Use the format
<DiskName>=<DiskSignature>.
For example,
Harddisk3=0x1234
Harddisk4=0xABCD.

[<<DiskName>=<DiskSignature>>...]

Format for indicating the disk name and signature. See the example below.

Example

```
vxabr write_signature Harddisk3=0x1234 Harddisk4=0xABCD
```

Manually adds the 0x1234 signature to Harddisk3 and 0xABCD signature to Harddisk4.

vxsnap

The `vxsnap` CLI utility lets you perform multiple snapshots at the same time. To use `vxsnap`, you must have InfoScale Enterprise installed.

Note: FlashSnap is available only with the Enterprise edition. See the "Licensing" section in the *Veritas InfoScale Installation and Upgrade Guide* for details.

Note: Dynamic disks belonging to a Microsoft Disk Management Disk Group do not support the `vxsnap` command.

The `vxsnap` CLI utility makes use of both FlashSnap and Volume Shadow Copy Service (VSS) technology to create high-quality snapshots that can be done when application files are open. VSS can quiesce the application for the moment when the snapshot is created and then resume the application immediately after the snapshot; but a VSS-aware application must be used, such as Microsoft SQL Server. The `vxsnap` CLI utility acts as a VSS requestor, which lets you use SFW to initiate VSS snapshots at any time.

See "[Volume Shadow Copy Service \(VSS\)](#)" on page 289.

The `vxsnap` utility is designed to create and restore multiple snapshot sets.

The `vxsnap` CLI utility lets you name the snapshot volume. You must create the mirror for the volume before using the `vxsnap` CLI utility. The `vxsnap prepare` command can be used to create the mirror for the volume.

The `vxsnap` CLI utility allows a traditional full snapshot or a Microsoft Copy on Write (COW) snapshot. Microsoft COW snapshots have the benefit of being faster than full snapshots and take up less storage space.

A COW snapshot creates a shadow copy that is differential. When a change to the original volume occurs, the block that is about to be modified is read and then written to a "differences area", which preserves a copy of the data block before it is overwritten with the change. Thus making a copy of the block before it is changed. Using the blocks in the differences area and unchanged blocks in the original volume, a copy of the volume can be constructed that represents the volume in the state before any changes occurred.

The benefit of the COW snapshot is that it creates shadow copies very rapidly because it only writes the changed blocks at a given time. Using the blocks in the differences area and unchanged blocks in the original volume, the initial state of the original volume can be reconstructed.

Note: Either the `prepare` or `start` keyword may be used in the CLI, however `prepare` is recommended.

Note: For `vxsnap` examples, see the application-specific Solutions guide. For example, for Microsoft SQL Server, refer to the *Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*.

Note: To use `vxsnap`, the disk group must be a Storage Foundation for Windows 4.0 or later version. You must upgrade the disk groups that are created with an earlier version of Volume Manager for Windows before using the `vxsnap` utility.

Note: The `vxsnap` commands must be invoked on a local system.

<code>prepare</code>	Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes Note: Either the <code>prepare</code> or <code>start</code> keyword may be used in the CLI, however <code>prepare</code> is recommended.
<code>create</code>	Creates simultaneous snapshots of all volumes in the specified component, providing a point-in-time snapshot set.
<code>reattach</code>	Reattaches and desynchronizes an existing snapshot set to the original database volumes.
<code>restore</code>	Restores a volume, database, or component from a snapshot set.
<code>refresh</code>	Refreshes all VSS Writers and component objects.
<code>diffarea</code>	Manages shadow storage area for COW snapshots. For example, the management of shadow storage area creation, size, and removal.
<code>delete</code>	Deletes COW snapshots corresponding to the information in the specified snapshot set.

Type the following sequence to view a description of the command syntax:

```
vxsnap <Keyword> -?
```

vxsnap prepare

For Exchange and volume operations:

```
vxsnap prepare component=<ComponentName>/writer=<WriterName>  
[-b] [source=<Volume>/harddisk=<Hardisk>...]
```

Creates snapshot mirrors of the volumes in the specified component and eliminates the need for multiple `vxassist prepare` commands. The snapshot mirrors remain attached to and synchronized with the original volumes.

Note: Either the `prepare` or `start` keyword may be used in the CLI, however `prepare` is recommended.

The following attributes apply:

<code>component=<ComponentName></code>	Name of the component; for Exchange, this is the storage group name found in the Exchange System Manager, for example, "First Storage Group".
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer".
<code>-b</code>	Run the process as a background process.
<code>source=<Volume></code>	Indicates the source volume for the snapshot mirror that is specified by a drive letter, drive path (mount point), or volume name of the form " <code>\\?\Volume{GUID}</code> ".
<code>harddisk=<Harddisk></code>	Name of the disk where the snapshot mirror is created, for example, <code>harddisk2</code> .

Example

```
vxsnap prepare component=SG2/writer="MicrosoftExchange Writer"  
source=L:/harddisk=harddisk2source=M:/harddisk=harddisk3
```

This command creates snapshot mirrors of all the volumes that the Exchange component contains. The snapshot mirror of the volume that is mounted on L: is created on disk 2 and the snapshot mirror of the volume that is mounted on M: is created on disk 3.

For Enterprise Vault operations:

```
vxsnap prepare [-b] writer=<WriterName>  
[site=<SiteName>[/VSG=<VSGName>[/VS=<VSName>]]]  
[[/]component=<ComponentName>]
```

Creates snapshot mirrors of the volumes in the specified site, vault store group, vault store, or component. The snapshot mirrors remain attached to and synchronized with the original volumes.

Note: Either the `prepare` or `start` keyword may be used in the CLI, however `prepare` is recommended.

The following attributes apply:

<code>-b</code>	Run the process as a background process.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, EnterpriseVault. Required if you specify the component.
<code>site=<SiteName></code>	Name of the Enterprise Vault Site.
<code>VSG=<VSGName></code>	Name of the Enterprise Vault Vault Store Group
<code>VS=<VSName></code>	Name of the Enterprise Vault Vault Store
<code>component=<ComponentName></code>	Name of the Enterprise Vault component. For example, VaultStore database, Fingerprint database, or Volume Component, such as index, partitions, etc.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Additional information about the `prepare` operation for Enterprise Vault:

- Specifying only the writer attribute (as in `vxsnap prepare writer=EnterpriseVault`), results in the preparation of the components of the entire Enterprise Vault farm. This would include the Enterprise Vault directory database, monitoring database, audit database, and reporting database.
- Specifying the site attribute (as in `vxsnap prepare writer=EnterpriseVault Site=Site1`), results in the preparation of the components at the Enterprise Vault site level. This prepares the directory database and all the components (SQL and volume) in the specified site.
- Specifying the vault store group attribute (as in `vxsnap prepare writer=EnterpriseVault Site=Site1/VSG=VSG1`), results in the preparation of the fingerprint database and all the components in all vault stores in the specified vault storage group.
- Specifying the vault store attribute (as in `vxsnap prepare writer=EnterpriseVault Site=Site1/VSG=VSG1/VS=VS1`), results in the preparation of the vault store database and all the partitions in the specified vault store.

- **Specifying the component attribute (as in `vxsnap prepare writer=EnterpriseVault Site=Site1/VSG=VSG1/VS=VS1 Component=<Component Name>`), results in the preparation of the specified component. The component can be an Enterprise Vault SQL database or a volume component (e.g.vault store database, fingerprint database, directory database, monitoring database, reporting database, audit database, partition, index volume, etc.).**
 - For vault store components: When the vault store database component is specified, the open partition in the vault store is also prepared.
 - For vault store components: When the open partition component is specified, the vault store database in the vault store is also prepared.
 - For different components: When different components reside on a common volume, the volume is prepared only once.
- Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, `%VMPATH%\logs\EVstatus.log`. The log contains information about the success or failure of the operation for the components. In the event that the prepare of a volume for a component fails, the operation continues to prepare the remaining volumes of the component and any other requested components. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

Example

```
vxsnap prepare writer=ev site=site1/vsg=vsg1/vs=vs1
site=site2 component="Index Location (V:)"
component="Directory DB"
```

The command prepares all the volumes for a snapshot for the Enterprise Vault Store vs1 of site1/vsg1, Index Location (V:) of site2 and Directory DB.

For SQL operations:

```
vxsnap prepare component=<ComponentName>/writer=<WriterName>
[server=<ServerName>] [instance=<InstanceName>] [-b]
[source=<Volume>/harddisk=<Harddisk>...]
```

Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes.

Note: Either the `prepare` or `start` keyword may be used in the CLI, however `prepare` is recommended.

The following attributes apply:

<code>component=<ComponentName></code>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.
<code>server=ServerName</code>	SQL Server server name.
<code>instance=InstanceName</code>	SQL Server instance name.
<code>-b</code>	Run the process as a background process.
<code>source=<Volume></code>	Indicates the source volume for the snapshot mirror that is specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\".
<code>harddisk=<Harddisk></code>	Name of the disk where the snapshot mirror is created, for example, harddisk2.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Example

```
vxsnap prepare component=billing_DB/writer=SQLServerWriter  
source=L:/harddisk=harddisk3  
source=M:/harddisk=harddisk3
```

This command creates snapshot mirrors of all the volumes that the billing_DB database contains. The snapshot mirror of the volume mounted on L: will be created on harddisk 3 and the snapshot mirror of the volume mounted on M: is also created on harddisk 3.

Note: Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each node. This is applicable for SQL Server 2012 or later.

vxsnap create

For Exchange operations:


```
vxsnap -x <Filename> create [-g <dynamicdiskgroupname>]  
source=<Volume>[/DriveLetter=<DriveLetter>] [/DrivePath=<DrivePath>]  
[/Newvol=<NewVolName>] [/Plex=<PlexName>] ...  
[writer=<WriterName>] [component=<ComponentName>]  
[backuptype=<Backuptype>] [-e] [-o] [-c] [secHosts=<SecondaryHosts>]
```

Creates snapshot(s) of the specified volume(s) or component. Allows the volumes to be snapshotted simultaneously.

Separate source volumes and attributes with forward slashes, not spaces. Source and snapshot volume attributes are paired. You must specify the source volume if you choose to specify the snapshot volume plex, drive letter, drive path, or volume name.

The following attributes apply:

<code>-x <Filename></code>	Indicates the name to be assigned to the XML metadata file that is created with the command. The file name must include the ".xml" extension. To specify a location other than the default location for the file, you must enter the full path for the file (e.g. J:\XML\Image1.xml).
<code>-g <DynamicDiskGroupName></code>	Name of the disk group that contains the source volume. Required when specifying volume name.
<code>source=<Volume></code>	Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified component (for example, Exchange storage group).
<code>[/plex=<PlexName>]</code>	Specifies the name of the mirror or plex that is to be detached. Use this parameter if there are multiple snap plexes available to be snapshotted.
<code>[/DriveLetter=<DriveLetter>]</code>	The drive letter to be assigned to the new snapshot volume.

[/DrivePath=<DrivePath>]	The drive path to be assigned to the new snapshot volume. The drive path must reference an empty local NTFS folder, which was created beforehand. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.
[/Newvol=<NewVolName>]	Specifies the name of the new snapshot volume that is created. If the name is not specified, then the default snapshot volume name begins with "SnapV" followed by a 12 character alphanumeric string. The 12 character alphanumeric string is a coded hexadecimal timestamp based upon the time that the snapshot was taken. For example, a default Snap Volume Name would be "SnapV4A7337AE038A". The full device path becomes: \Device\HarddiskDmVolumes \ <diskgroupname>\<newvolname>< td=""></diskgroupname>\<newvolname><>
writer=<WriterName>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or "Microsoft Exchange Writer Replica". If you plan to be able to do a roll-forward recovery to the point of failure, you must specify the writer.
component=<ComponentName>	Name of the component; for Exchange, this is the storage group name found in the Exchange System Manager, for example, "First Storage Group". If you plan to be able to do a roll-forward recovery to the point of failure, you must use this component.
backuptype=<Backuptype>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default. Copy backup creates a copy of the database and transaction logs volumes. Full backup creates a copy of the database and transaction logs volumes, runs Eseutil to check for consistency, and if consistent, truncates the transaction logs.

<code>-e</code>	Runs the Eseutil consistency check for database and log files. Eseutil is run automatically with a full backup, but must be optionally specified for a copy backup.
<code>-o</code>	Allows an existing XML file of the same name to be overwritten. If <code>-o</code> is not specified the <code>vxsnap create</code> command does not overwrite an existing XML file of the same name and the operation fails.
<code>-c</code>	Specifies a COW snapshot. In preparation for a COW snapshot, the <code>vxsnap diffarea</code> command must be used to create the shadow storage volume. The following <code>vxsnap</code> attributes are not compatible with COW snapshots and result in an error: <ul style="list-style-type: none">■ <code>/DriveLetter=<driveLetter></code>■ <code>/DrivePath=<drivePath></code>■ <code>/Newvol=<newVolName></code>■ <code>/Plex=<plexName></code>■ <code>secHosts=<secondary hosts></code>
<code>sechosts=<SecondaryHosts></code>	Applies to Volume Replicator (Volume Replicator) environment only. Comma-separated list of secondary hosts on which a synchronized snapshot is to be taken.

Note: Any text string that contains spaces must be enclosed in quotation marks.

This example creates a snapshot of the replica of the component mailbox. It specifies that the VSS Writer, Microsoft Exchange Writer Replica, is used for the snapshot.

```
vxsnap -x snapdata.xml create
-c writer="Microsoft Exchange Writer"
component=mb1 backuptype=COPY -E
```

Creates a COW snapshot of component mb1 of Exchange server. A copy backup is created and an Eseutil consistency check is performed on the snapshot volumes.

For Enterprise Vault operations:

```
vxsnap -x <Filename> create writer=<WriterName>  
[site=<SiteName>[/VSG=<VSGName>[VS=<VSName>]]]  
[[/]/component=<ComponentName>[backuptype=FULL|COPY] [-o]
```

Creates snapshot(s) of the Enterprise Vault components. Allows the volumes to be snapshotted simultaneously.

The following attributes apply:

<code>-x<Filename></code>	Indicates the name to be assigned to the XML metadata file that is created with the command. The file name must include the ".xml" extension. To specify a location other than the default location for the file, you must enter the full path for the file (e.g. J:\XML\Image1.xml).
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, EnterpriseVault. Required if you specify the component.
<code>site=<SiteName></code>	Name of the Enterprise Vault site.
<code>VSG=<VSGName></code>	Name of the Enterprise Vault Vault Store Group.
<code>VS=<VSName></code>	Name of the Enterprise Vault Vault Store.
<code>component=<ComponentName></code>	Name of the Enterprise vault component. For example, Vault Store database, Fingerprint database, or Volume component, such as index, partitions, etc.
<code>backuptype=FULL COPY</code>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default.
<code>-o</code>	Allows an existing XML file of the same name to be overwritten. If <code>-o</code> is not specified the <code>vxsnap create</code> command does not overwrite an existing XML file of the same name and the operation fails.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Additional information about the create operation for Enterprise Vault:

- Specifying only the writer attribute (as in `vxsnap -x xmlfile create writer=EnterpriseVault`), results in the snapshot of the components of the entire Enterprise Vault farm. This would include the Enterprise Vault directory database, monitoring database, audit database, and reporting database.
- Specifying the site attribute (as in `vxsnap -x xmlfile create writer=EnterpriseVault Site=Site1`), results in the snapshot of the components at the Enterprise Vault site level. This prepares the directory database and all the components (SQL and volume) in the specified site.
- Specifying the vault store group attribute (as in `vxsnap -x xmlfile create writer=EnterpriseVault Site=Site1/VSG=VSG1`), results in the snapshot of the fingerprint database and all the components in all vault stores in the specified vault storage group.
- Specifying the vault store attribute (as in `vxsnap -x xmlfile create writer=EnterpriseVault Site=Site1/VSG=VSG1/VS=VS1`), results in the snapshot of the vault store database and all the partitions in the specified vault store.
- Specifying the component attribute (as in `vxsnap -x xmlfile create writer=EnterpriseVault Site=Site1/VSG=VSG1/VS=VS1 Component=<Component Name>`), results in the snapshot of the specified component. The component can be an Enterprise Vault SQL database or a volume component (e.g.vault store database, fingerprint database, directory database, monitoring database, reporting database, audit database, partition, index volume, etc.).
 - For vault store components: When the vault store database component is specified, the open partition in the vault store is also snapshotted.
 - For vault store components: When the open partition component is specified, the vault store database in the vault store is also snapshotted.
 - For different components: When different components reside on a common volume, the volume is snapshotted only once.
- Upon completion of the operation, the status (successor failure) of the selected components is recorded in a log, `%VMPATH%\logs\EVstatus.log`. The log contains information about the success or failure of the operation for the components. In the event that the snapshot of a volume for a component fails, the operation does not snapshot the remaining volumes of the component nor any other requested components that are remaining. If the operation succeeds for all the volumes of a component, then the status of the component is logged as SUCCESS. If the operation fails for any one of the volumes of the component, then the status of all the components of that EnterpriseVault server is logged as FAILED along with the cause of the failure. The status of all components that

had been already snapshotted earlier before the failed components are logged as SUCCESS, and the status of the remaining components that had not been snapshotted are logged as SKIPPED.

Examples

```
vxsnap -x snapdata.xml create writer=EnterpriseVault
site=site1/vsg=vsg1/vs=vs1
site=site2 component="Index Location (V:)" component="Directory DB"
```

This example makes a snapshot of vault store "vs1" of site1/vsg1, "Index Location (V:)" of site2 and component "Directory DB".

```
vxsnap -x snapdata.xml create writer=EnterpriseVault site=site1
```

This example makes a complete snapshot of site1 (including the directory database).

For SQL operations:

```
vxsnap -x <Filename> create [-g <dynamicdiskgroupname>]
source=<Volume>[/DriveLetter=<DriveLetter>]/[DrivePath=<DrivePath>]
[/Newvol=<NewVolName>]/[Plex=<PlexName>]...
[writer=<WriterName>] [component=<ComponentName>]
[server=<ServerName>] [instance=<InstanceName>]
[backuptype=FULL|COPY] [-o] [-c] [secHosts=<SecondaryHosts>]
```

Creates snapshot(s) of the specified volume(s) or SQL database. Allows the volumes to be snapshotted simultaneously.

Separate source volumes and attributes with forward slashes, not spaces. Source and snapshot volume attributes are paired. You must specify the source volume if you choose to specify the snapshot volume plex, drive letter, drive path, label, or volume name.

The following attributes apply:

- | | |
|---------------------------|---|
| -x <Filename> | Indicates the name to be assigned to the XML metadata file that is created with the command. The file name must include the ".xml" extension. To specify a location other than the default location for the file, you must enter the full path for the file (e.g. J:\XML\Image1.xml). |
| -g <DynamicDiskGroupName> | Name of disk group that contains the source volume. Required when specifying volume name. |

<code>source=<Volume></code>	Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified component.
<code>[/plex=<PlexName>]</code>	Specifies the name of the mirror or plex that is to be detached. Use this parameter if there are multiple snapplexes available to be snapshotted.
<code>[/DriveLetter=<DriveLetter>]</code>	The drive letter to be assigned to the new snapshot volume.
<code>[/DrivePath=<DrivePath>]</code>	The drive path to be assigned to the new snapshot volume. The drive path must reference an empty local NTFS folder, which was created beforehand. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.
<code>[/Newvol=<NewVolName>]</code>	<p>Specifies the name of the new snapshot volume that is created.</p> <p>If the name is not specified, then the default snapshot volume name begins with "SnapV" followed by a 12 character alphanumeric string. The 12 character alphanumeric string is a coded hexadecimal timestamp based upon the time that the snapshot was taken. For example, a default Snap Volume Name would be "SnapV4A7337AE038A".</p> <p>The full device path becomes:</p> <pre>\Device\HarddiskDmVolumes <DiskGroupName>\<NewVolName></pre>
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer. Required if you specify the component.

<code>component=<ComponentName></code>	Name of the database. The command prepares mirrors for both the database and log volumes of the SQL database.
<code>server=ServerName</code>	SQL Server server name.
<code>instance=InstanceName</code>	SQL Server instance name.
<code>backuptype=FULL COPY</code>	Specifies the type of backup, either a Full or Copy. If no option is specified then Copy is the default. To back up logs in SQL Server so that you can restore the database using SQL log replay, at least one Full backup must have been created earlier.
<code>-o</code>	Allows an existing XML file of the same name to be overwritten. If <code>-o</code> is not specified the <code>vxsnap create</code> command does not overwrite an existing XML file of the same name and the operation fails.
<code>-c</code>	Specifies a COW snapshot. In preparation for a COW snapshot, the <code>vxsnap diffarea</code> command must be used to create the shadow storage volume. The following <code>vxsnap</code> attributes are not compatible with COW snapshots and result in an error: <ul style="list-style-type: none">■ <code>/DriveLetter=<driveLetter></code>■ <code>/DrivePath=<drivePath></code>■ <code>/Newvol=<newVolName></code>■ <code>/Plex=<plexName></code>■ <code>secHosts=<secondary hosts></code>
<code>sechosts=<SecondaryHosts></code>	Applies to Volume Replicator (Volume Replicator) environment only. Comma-separated list of secondary hosts on which a synchronized snapshot is to be taken.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Examples


```
vxsnap -x billing_DB.xml create
source=G:/Newvol=billing_data
source=H:/Newvol=billing_log
writer=SQLServerWriter component=billing_DB
backuptype=full
```

This example creates a snapshot set based on the component `billing_DB` which contains volume `G`, the database volume, and volume `H`, the log volume. The snapshot volumes are named `billing_data` and `billing_log`, respectively. The XML file, `billing_DB.xml`, is used to store the VSS metadata that identifies the snapshot set. This file is used in the command to reattach the snapshot set to resynchronize it and in the command to recover the database using the snapshot set.

```
vxsnap -x snapdata.xml create
-C writer=sqlserverwriter
component=testDB server=localhost
```

Creates a COW snapshot of SQL database `testDB` used in the default instance of the `localhost`.

For volume operations:

```
vxsnap -x <Filename> create [-g <dynamicdiskgroupname>]
source=<Volume>[/DriveLetter=<DriveLetter>][/DrivePath=<DrivePath>]
[/Newvol=<NewVolName>][/Plex=<PlexName>]...[-o] [-c]
```

Creates snapshot(s) of the specified volume(s). Allows volumes to be snapshotted simultaneously.

Separate source volumes and attributes with forward slashes, not spaces. Source and snapshot volume attributes are paired. You must specify the source volume if you choose to specify the snapshot volume plex, drive letter, drive path, label, or volume name.

The following attributes apply:

<code>-x <Filename></code>	Indicates the name to be assigned to the XML metadata file that is created with the command. The file name must include the ".xml" extension.
<code>-g <DynamicDiskGroupName></code>	Name of disk group that contains the source volume. Required when specifying volume name.

`source=<Volume>`

Indicates the source volume for the split-mirror snapshot specified by a drive letter, drive path (mount point), or volume name of the form "\\?\Volume{GUID}\". Repeat this parameter for each volume associated with the specified component.

`[/plex=<PlexName>]`

Specifies the name of the mirror or plex that is to be detached. Use this parameter if there are multiple snapplexes available to be snapshotted.

`[/DriveLetter=<DriveLetter>]`

The drive letter to be assigned to the new snapshot volume.

`[/DrivePath=<DrivePath>]`

The drive path to be assigned to the new snapshot volume. The drive path must reference an empty local NTFS folder, which was created beforehand. The path must include the drive letter and folder to be mounted, for example, C:\DB1VOL.

`[/Newvol=<NewVolName>]`

Specifies the name of the new snapshot volume that is created.

If the name is not specified, then the default snapshot volume name begins with "SnapV" followed by a 12 character alphanumeric string. The 12 character alphanumeric string is a coded hexadecimal timestamp based upon the time that the snapshot was taken. For example, a default Snap Volume Name would be "SnapV4A7337AE038A".

The full device path becomes:

`\Device\HarddiskDmVolumes`

`\<DiskGroupName>\<NewVolName>`

`-o`

Allows an existing XML file of the same name to be overwritten. If `-o` is not specified the `vxsnap create` command does not overwrite an existing XML file of the same name and the operation fails.

-c

Specifies a COW snapshot.

In preparation for a COW snapshot, the vxsnap diffarea command must be used to create the shadow storage volume. The following vxsnap attributes are not compatible with COW snapshots and result in an error:

- /DriveLetter=<driveLetter>
- /DrivePath=<drivePath>
- /Newvol=<newVolName>
- /Plex=<plexName>
- secHosts=<secondary hosts>

Note: Any text string that contains spaces must be enclosed in quotation marks.

Examples

```
vxsnap -x snapdata.xml create
source=E:\Data\DB1/DrivePath=E:\Backup\DB1
source=E:\Data\DB2/DrivePath=E:\Backup\DB2
```

This command does not specify a specific storage group (component) but rather creates snapshots from the volumes mounted on E:\Data\DB1 and E:\Data\DB2. The resulting snapshot volumes are assigned mount points E:\Backup\DB1 and E:\Backup\DB2, respectively. The metadata involved in this operation is stored in snapdata.xml.

```
vxsnap -x snapdata.xml create
-C source=D:
```

Creates a COW snapshot of volume D:

vxsnap reattach

For Exchange operations:

```
vxsnap -x <Filename> [-f] [-b] reattach[writer=<WriterName>]
[secHosts=<SecondaryHosts>]
```

This command reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

The following attributes apply:

<code>-x <Filename></code>	<p>The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file.</p> <p>Note: This file is deleted after the reattach operation has completed successfully.</p>
<code>-f</code>	<p>Forces the reattach. Make sure that the volume is not in use by another application before using this command. Use this option with care.</p>
<code>-b</code>	<p>Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.</p>
<code>writer=<WriterName></code>	<p>Unique ID of the VSS writer, for example, "Microsoft Exchange Writer".</p>
<code>sechosts=<SecondaryHosts></code>	<p>Applies to Volume Replicator (Volume Replicator) environment only. Comma-separated list of secondary hosts on which a synchronized snapshot is to be taken.</p>

Note: Make sure that the snapshot volumes are not in use before using this command.

Example

```
vxsnap -x snapdata.xml reattach
```

This command uses the information in the `snapdata.xml` file to reattach and resynchronize all the volumes in the snapshot set. This xml file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the `vxsnap create` command is issued.

For Enterprise Vault operations:

```
vxsnap -x <Filename> [-f] [-b] reattach [writer=<WriterName>  
[site=<SiteName>[/VSG=<VSGName>[VS=<VSName>]]]  
[[/]component=<ComponentName>
```

This command reattaches and resynchronizes the snapshot volumes in the snapshot set to the original Enterprise Vault component volumes.

The following attributes apply:

<code>-x<Filename></code>	The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file. Note: This file is deleted after the reattach operation has completed successfully.
<code>-f</code>	Forces the reattach. Make sure that the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, EnterpriseVault or the GUID for the writer.
<code>site=<SiteName></code>	Name of the Enterprise Vault Site.
<code>VSG=<VSGName></code>	Name of the Enterprise Vault Vault Store Group.
<code>VS=<VSName></code>	Name of the Enterprise Vault Vault Store.
<code>component=<ComponentName></code>	Name of the Enterprise vault component. For example, Vault Store database, Fingerprint database, or Volume component, such as index, partitions, etc.

Note: Any text string that contains spaces must be enclosed in quotation marks.

Note: Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, `%VMPATH%\logs\EVStatus.log`. The log contains information about the success or failure of the operation for the components. In the event that the reattach of a volume for a component fails, the operation continues to reattach the remaining volumes of the component and any other requested components. The components that successfully complete the operation are removed from the snapshot set. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

Example

```
vxsnap -x snapdata.xml reattach writer=EnterpriseVault
site=site1/vsg=vsg1/vs=vs1 site=site2
component="Index Location (V:)"
component="Directory DB"
```

This example reattaches all the snapshot volumes for the vault store vs1 of site1/vsg1, "Index Location (V:)" of site2 and "Directory DB".

For SQL operations:

```
vxsnap -x <Filename> [-f] [-b] reattach[writer=<WriterName>]
[secHosts=<SecondaryHosts>]
```

This command reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

The following attributes apply:

<code>-x <Filename></code>	The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file. Note: This file is deleted after the reattach operation has completed successfully.
<code>-f</code>	Forces the reattach. Make sure that the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, SQLServerWriter or the GUID for the writer.
<code>sechosts=<SecondaryHosts></code>	Applies to Volume Replicator (Volume Replicator) environment only. Comma-separated list of secondary hosts on which a synchronized snapshot is to be taken.

Note: Make sure that the snapshot volumes are not in use before using this command.

Example

```
vxsnap -x billing_DB.xml reattach writer=SQLServerWriter
```

This command uses the information in the snapdata.xml file to reattach and resynchronize all the volumes in the snapshot set. This xml file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the `vxsnap create` command is issued.

Note: Any text string that contains spaces must be enclosed in quotation marks.

For volume operations:

```
vxsnap -x <Filename> [-f] [-b] reattach
```

This command reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

The following attributes apply:

- x <Filename> The file that is created by the `vxsnap create` command. Each snapshot set must have a unique name for the metadata file.
Note: This file is deleted after the reattach operation has completed successfully.
- f Forces the reattach. Make sure that the volume is not in use by another application before using this command. Use this option with care.
- b Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.

Note: Make sure that the snapshot volumes are not in use before using this command.

Example

```
vxsnap -x snapdata.xml reattach
```

This command uses the information in the snapdata.xml file to reattach and resynchronize all the volumes in the snapshot set. This xml file is deleted after the reattach operation has completed successfully. The snapshot volumes remain synchronized with the original volumes until the `vxsnap create` command is issued.

vxsnap restore

For Exchange operations:

```
vxsnap -x <Filename> [-f] [-b] [-r] [-a] restore  
restoreType=<PIT|POF> writer=WriterName  
[subComponent=<subComponentName>]
```

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to restore data, for example, after an original volume has become corrupted. You can restore the data either to the point in time that the snapshot set was last refreshed or to the point of failure of a single database.

(COW snapshots can be used with this command.)

Note: For Exchange 2010, Restore to Recovery Database is not supported in SFW.

Note: After completing a point of failure (POF) recovery of a single database, Veritas recommends using the `vxsnap reattach` command to reattach and resynchronize the other databases in the Exchange DAG and to use the `vxsnap create` command to create a new snapshot set.

Implementing the point of failure recovery requires that the `writer=WriterName` and the `component=<ComponentName>` parameters were specified when the snapshot set was created.

The following attributes apply:

<code>-x <Filename></code>	The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file. When the full path for the <code><Filename></code> is not specified, the <code>writer=<WriterName></code> attribute is required.
<code>-f</code>	Forces the snapback. Make sure that the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.

<code>-r</code>	<p>Recover one or more of the original missing volumes. The following example shows additional required steps.</p> <p>This option cannot be specified to recover using a COW snapshot.</p>
<code>-a</code>	<p>Dismount the databases before the restore operation and then mount the database after the restore operation.</p>
<code>restoreType=<PIT POF></code>	<p>PIT specifies a restore to the point in time that the snapshot set was created or last refreshed.</p> <p>POF specifies a roll-forward recovery to the point of failure.</p>
<code>writer=WriterName</code>	<p>The name for the Exchange Server VSS Writer; used to locate the default directory to search for the XML metadata file.</p>
<code>subComponent=<subComponentName></code>	<p>Name of the subcomponent to be restored. In Exchange, a subcomponent is a mailbox store (database). Use this attribute only in a point of failure recovery.</p> <p>Note: The subcomponent attribute is not supported for Exchange 2010.</p>

Note: Before using this command, make sure that the source volumes and the snapshot volumes are not in use. Use the `[-a]` attribute to dismount and mount the databases automatically or use the Exchange System Manager to dismount all the databases in the DAG and then mount them after the command is completed.

Examples

```
vxsnap -x snapdata.xml restore  
writer="Microsoft Exchange Writer"
```

This command uses the information in the `snapdata.xml` file to restore all the volumes in the snapshot set identified in that file to the point in time the snapshot set was created or last refreshed.

Point-in-time recovery

```
vxsnap -x snapdata.xml restore restoreType=PIT  
writer="Microsoft Exchange Writer"
```

This command uses the information in the `snapdata.xml` file to restore all the volumes in the snapshot set identified in that file to the point in time the snapshot set was created or last refreshed.

Roll-Forward Recovery to the Point of Failure

```
vxsnap -x snapdata.xml restore restoreType=POF  
writer="Microsoft Exchange Writer"
```

This command uses the information about the storage group that is specified in the `snapdata.xml` file to snapback the database volumes and then use current transaction logs to roll forward to the point of failure.

Roll-Forward Recovery to the Point of Failure of a Single Database

```
vxsnap -x snapdata.xml restore restoreType=POF  
writer="Microsoft Exchange Writer" subcomponent=DB1
```

This command restores the specified database (subcomponent) DB1 and then uses current transaction logs to roll forward only that database to the point of failure.

Recovery After Hardware Failure

You can use the `-r` switch to perform a VSS-integrated recovery after a hardware failure. The following recovery scenarios are possible if the complete snapshot set including the XML metadata file is available.

Note: For more information about the `-r` switch, see the Snapshot Solutions section in the *Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange*, which is included with the software.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume that is associated with each missing production volume. Note the drive letter or mount point of each volume.
- Use Exchange System Manager to dismount all remaining databases in the storage group.
- Delete the missing volumes from Storage Foundation for Windows.

- Replace the failed hardware and add the new disks to the dynamic disk group.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
- Perform a VSS-integrated recovery by including the `-r` switch in the `vxsnap restore` command. For example, type:

```
vxsnap -x snapdata.xml -r restore restoreType=PIT
writer="Microsoft Exchange Writer"
```

This command uses the information the `snapdata.xml` file to restore all the volumes in the snapshot set identified in that file to the point in time the snapshot set was created or last refreshed.

Warning: Before using the `vxsnap restore` command, verify that you have correctly assigned the drive or mount point to each volume and that you have accounted for all the volumes in the storage group (component).

For Enterprise Vault operations

```
vxsnap -x <Filename> [-b] [-f] [-r] [-a] restore
  writer=WriterName [site=<siteName>[/VSG=<VSGName>[/VS=<VSName>]]]
  [[/]component=<ComponentName>] {RestoreType=[RECOVERY|NO_RECOVERY]}
```

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to recover a corrupted or missing Enterprise Vault component.

Exclusive access to the Enterprise Vault component is required for this operation.

Before using this command verify that the source volumes and the snapshot volumes are not in use.

The following attributes apply:

<code>-x <Filename></code>	The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>-f</code>	Forces the snapback. Use this option with care.

<code>-r</code>	Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume.
<code>-a</code>	Dismount the databases before the restore operation and then mount the database after the restore operation.
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, EnterpriseVault or the GUID for the writer.
<code>site=<SiteName></code>	Name of the Enterprise Vault site.
<code>VSG=<VSGName></code>	Name of the Enterprise Vault Store Group.
<code>VS=<VSName></code>	Name of the Enterprise Vault Store.
<code>component=<ComponentName></code>	Name of the Enterprise vault component. For example, Vault Store database, Fingerprint database, or Volume component, such as index, partitions, etc.
<code>restoreType= [RECOVERY NO_RECOVERY]</code>	<p>Specifies the type of database recovery, either recovery, or no recovery:</p> <p>With <code>RECOVERY</code> database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.</p> <p>To back up logs so that you can restore the database using log replay, at least one Full backup must have been created earlier.</p> <p><code>NO_RECOVERY</code> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in a loading state so that you can manually replay backup logs to a specific point in time.</p>

Note: Upon completion of the operation, the status (success or failure) of the selected components is recorded in a log, %VMPATH%\logs\EVStatus.log. The log contains information about the success or failure of the operation for the components. In the event that the restore of a volume for a component fails, the operation continues to restore the remaining volumes of the component and any other requested components. The components that successfully complete the operation are removed from the snapshot set. If the operation succeeds for all the volumes of a component, then the status of the component is logged as a success. If the operation fails for any one of the volumes of the component, then the status of the component is logged as a failure along with the cause of failure.

The following is an example of the command.

```
vxsnap -x snapdata.xml -a restore RestoreType=RECOVERY  
writer=EnterpriseVault  
site=sitel/vsg=vsg1/vs=vs1 component="Directory DB"
```

This command restores the vault store vs1 and component Directory DB using the metadata from snapdata.xml.

For SQL operations:

```
vxsnap -x <Filename> [-f] [-b] [-r] restore  
[restoreType=[RECOVERY|NO_RECOVERY]]  
[noLogs|logFiles=<tlog1,tlog2,...>] writer=WriterName
```

Uses the snapshot volumes in a snapshot set created by the `vxsnap create` command to recover a corrupted or missing SQL Server database. Exclusive access to the SQL Server database is required for this operation.

(COW snapshots can be used with this command.)

Before using this command verify that the source volumes and the snapshot volumes are not in use.

The following attributes apply:

`-x <Filename>`

The file that is created by the `vxsnap create` command. Each snapshot set must have a unique name for the metadata file.

When the full path for the `<Filename>` is not specified, the `writer=<WriterName>` attribute is required.

<code>-f</code>	Forces the snapback. Make sure that the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>-r</code>	<p>Recover even if original volume is not present. If this option is selected and the original volume is not present, the snapshot volume of the missing volume is changed from a read-only volume to a read-write volume.</p> <p>Use this option only with Recovery noLogs. After using this option you must explicitly assign the original drive letter/mount path of the missing volume to the snapshot volume in the VEA and then bring the database online.</p> <p>This option cannot be specified to recover using a COW snapshot.</p>
<code>restoreType=</code> <code>[RECOVERY NO_RECOVERY]</code>	<p>Specifies the type of database recovery, either recovery, or no recovery:</p> <p><code>RECOVERY</code> can be used with either the <code>noLogs</code> or <code>logFiles=tlog1,tlog2,...</code> attributes. <code>RECOVERY</code> leaves the database in an online state.</p> <p>To back up logs so that you can restore the database using log replay, at least one Full backup must have been created earlier.</p> <p><code>NO_RECOVERY</code> restores from the specified snapshot set to the time of the snapshot. No logs are applied and the database is left in a loading state so that you can manually replay backup logs to a specific point in time.</p>
<code>noLogs</code>	Database and transaction log files are restored from the snapshot set. No transaction backup logs are applied. The database is left in an operational state.
<code>logFiles=tlog1,tlog2,...</code>	Transaction log backup files to be applied with the <code>RECOVERY</code> option to achieve a point of failure recovery and leave the database in an online state. Each transaction log must have a unique name and be created using the "overwrite existing media" option.
<code>writer=WriterName</code>	The name for the SQL Server VSS Writer; used to locate the default directory to search for the XML metadata file. Specify <code>SQLServerWriter</code> .

The following are examples of the command:

- Recovering using snapshots without log replay

```
vxsnap -x TestDB.xml restore RestoreType=RECOVERY noLogs  
writer=SQLServerWriter
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and brings the database online. The database is restored to the time the snapshot set was created or last refreshed.

You can use the `-r` option with the `RECOVERY noLogs` restore type if a production volume is missing due to hardware failure:

```
vxsnap -x TestDB.xml -r restore RestoreType=RECOVERY noLogs  
writer=SQLServerWriter
```

- Recovering using snapshots and log replay

```
vxsnap -x TestDB.xml restore RestoreType=RECOVERY  
logFiles=c:\backup\tLog1.bak, c:\tLog2.bak  
writer=SQLServerWriter
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and then applies the specified transaction log backups (c:\backup\tLog1.bak and c:\tLog2.bak) and brings the database online.

- Restoring snapshots and manually applying logs

```
vxsnap -x TestDB.xml restore RestoreType=NO_RECOVERY  
writer=SQLServerWriter
```

This command uses the information in the TestDB.xml file to restore all the volumes in the snapshot set and leaves the database in a loading state so that backup logs can be manually restored to a specific point in time.

Note: For more information about the `-r` switch, see the *Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*.

For volume operations:

```
vxsnap -x <Filename> [-f] [-b] [-r] restore RestoreType=PIT  
[<Volumename|Driveletter|DrivePath> ...]
```

Uses the snapshots in a snapshot set created by the `vxsnap create` command to restore data, for example, after an original volume has become corrupted.

(COW snapshots can be used with this command.)

The following attributes apply:

<code>-x <Filename></code>	The file that is created by the <code>vxsnap create</code> command. Each snapshot set must have a unique name for the metadata file.
<code>-f</code>	Forces the snapback. Make sure that the volume is not in use by another application before using this command. Use this option with care.
<code>-b</code>	Resynchronizes the volume in the background. A new snapshot cannot be made until the resynchronization is complete.
<code>-r</code>	Recover one or more of the original volumes are missing. This option cannot be specified to recover using a COW snapshot.
<code>RestoreType=<PIT></code>	PIT specifies a restore to the point in time that the snapshot set was created or last refreshed.
<code>VolumeName</code>	Name of volume. For example, <code>\Device\HarddiskDmVolumes\DynamicGroup\Volume1</code> .
<code>DriveLetter</code>	Drive letter of the volume.
<code>DrivePath</code>	Drive path of the volume.

Examples

```
vxsnap -x snapdata.xml restore RestoreType=PIT
```

This command uses the information in the `snapdata.xml` file to restore all the volumes in the snapshot set to the point in time the snapshot set was created or last refreshed.

vxsnap refresh

For Exchange and SQL operations:

```
vxsnap refresh
```

This command refreshes all VSS Writers and component objects. (Exchange, SQL, Enterprise Vault, and COW operations)

Note: It is observed that when you add or remove the VSS components or when you connect to the VEA console for the first time, the updated VSS components are not displayed on the VEA console. To view the updated VSS components, you must manually refresh the VEA. To manually refresh the console, run either the `vxsnap refresh` command or perform a Refresh operation through the VEA console.

vxsnap diffarea

For COW snapshot operations:

```
Vxsnap diffarea <-a|-r|-c> source=<driveletter>  
/target=<driveletter>[/size=<size>]...
```

Manages creation, deletion, and modification of shadow storage area settings for COW snapshots.

The following attributes apply:

- a Create shadow storage area for volume that is specified in source. Shadow storage area resides on volume specified in target with specified size. The shadow storage area must reside in the same dynamic disk group as the source. It can reside on the same volume as the source volume.
- r Remove association of shadow storage area that is specified in target and volume that is specified in source.
- c Change the shadow storage area settings for an existing shadow storage area that is specified in the target and the volume that is specified in source.
- source Source volume for COW snapshot.
Note: A basic volume cannot be used as the source volume.
- target Volume for storing blocks of source volume before write operation.
Note: A basic volume cannot be used as the target volume.
Note: Target volume must be an NTFS volume.
- size Size of the volume that is shadow storage area.
If size is set to -1 or if size is not specified, then the size of the shadow storage area is unlimited.

Note: Use the VEA GUI to view the settings for shadow storage areas.

Examples

```
vxsnap diffarea -a source=E:/target=F:/size=500
source=G:/target=H:/size=300
```

This command creates two shadow storage areas. One is a 500MB shadow storage area on volume F that stores COW snapshots of volume E. The second is a 300MB shadow storage area on volume H that stores COW snapshots of volume G.

```
vxsnap diffarea -r source=E:/target=F:
```

This command removes the association between volume E and the shadow storage area on volume F.

vxsnap delete

For COW snapshot operations:

```
vxsnap -x <Filename> [-f] delete
[writer=WriterName]
```

Deletes the COW snapshots corresponding to the information in the specified XML file. The XML file is deleted after all the snapshots have been successfully deleted.

The following attributes apply:

<code>-x <Filename></code>	Name of the file that stores the metadata for COW snapshots. The file is created by the <code>vxsnap create</code> command.
<code>-f</code>	Forces the delete. Use this option with care. Use the <code>-f</code> option when one or more snapshots in the snapshot set do not exist. For example, use the <code>-f</code> option when one or more snapshots in the snapshot set no longer exist because Windows has deleted older snapshots to make room for newer snapshots.
<code>writer=WriterName</code>	The name for the VSS Writer that is associated with the snapshots to be deleted.

Examples

```
vxsnap -x snapdata.xml delete writer="Microsoft Exchange Writer"
```

Deletes the Exchange COW snapshots corresponding to information in `snapdata.xml`.

```
vxsnap -x snapdata.xml delete writer=sqlserverwriter
```

Deletes the SQL COW snapshots corresponding to information in `snapdata.xml`.

```
vxsnap -x snapdata.xml delete
```

Deletes the COW snapshots corresponding to information in snapdata.xml.

```
vxsnap -x cow1.xml -f delete
```

Forces the deletion of the COW snapshots corresponding to information in cow1.xml.

vxfsync

The vxfsync utility uses Fast File Resync to resynchronize a single file in a snapshotted volume to the original volume. It is designed to restore a good copy of one corrupt file in a database application that has volumes with multiple database files.

```
vxfsync -g <DynamicDiskGroupName> -m <Mastervolume> -s <Snapvolume>
-f <filename>
```

See [“Fast File Resync”](#) on page 287.

The vxfsync command is only available from the following Storage Foundation folder:

```
%ProgramFiles%\Veritas\Veritas Volume Manager
```

-g <DynamicDiskGroupName>	The name of the disk group containing the original and snapshotted volumes.
-m <Mastervolume>	The original volume.
-s <Snapvolume>	The name of the split-mirror snapshot volume.
-f <filename>	The file name of the file to be resynchronized to the original volume.

Example

```
vxfsync -g test -m vol1 -s vol1_snap01 -f test.dat
```

This command uses the snapshot volume (the replica), vol1_snap01, to resynchronize or restore the file test.dat on the master or original volume, vol1.

vxscrub

vxscrub is a command for removing unneeded entries in the registry that are a result of adding and deleting volumes. The command is designed to shrink the size of the registry by purging the system hive. Shrinking the size of the registry after

adding and deleting volumes avoids problems during system restart. You must restart the system after running `vxscrub`.

Warning: The `vxscrub` utility should be used only under the direction of Support. Use of this utility without supervision may cause performance problems when importing disk groups.

<code>[-a -b -l -v]</code>	Uninstalls all unused volumes, disks, or mounted devices only. Does not purge the system hive.
<code>-p</code>	Uninstalls all unused volumes, disks, and mounted devices, then purges the system hive.
<code>-forcepurge</code>	Purges the system hive only. Does not uninstall unused volumes, disks, or mounted devices.

Typing the following brings up a description of its syntax:

```
vxscrub
```

vxscrub [-a|-b|-l|-v]

```
vxscrub [-a|-b|-l|-v] [-d] [-m] [-q] [-c]
```

Uninstalls all unused volumes, disks, or mounted devices only. Does not purge the system hive.

Note: The system hive of the registry does not physically shrink using this command mode. Use the other command modes to physically shrink the system hive of the registry.

The following attributes apply:

<code>-a</code>	Uninstalls all unused volumes.
<code>-b</code>	Uninstalls all unused BASIC volumes.
<code>-l</code>	Uninstalls all unused LDM volumes.
<code>-v</code>	Uninstalls all unused SFW volumes.
<code>-d</code>	Uninstalls all unused disk instances.
<code>-m</code>	Cleans up the database of mounted devices.

- q Does not display any output.
- c Tests the uninstall and the cleanup operation. The uninstall and clean up operation is not performed.

Example

```
vxscrub -a -d -m
```

This command uninstalls all unused volumes, disks, and mounted devices.

vxscrub -p

```
vxscrub -p [-q]
```

Uninstalls all unused volumes, disks, and mounted devices, then the system hive is purged.

Note: You must restart the system after running this command.

The following attributes apply:

- p Uninstalls all unused volumes, disks, and mounted devices, then the system hive is purged.
- q Does not display any output.

Example

```
vxscrub -p
```

This command uninstalls all unused volumes, disks, and mounted devices; purges the system hive; and shrinks the size of the registry.

vxscrub -forcepurge

```
vxscrub -forcepurge [-c]
```

Purges the system hive only. Does not uninstall unused volumes, disks, or mounted devices.

Note: You must restart the system after running this command.

The following attributes apply:

- forcepurge Purges the system hive.
- c Tests the purge operation and reports the new size of the registry. The purge operation is not performed.

Example

```
vxscrub -forcepurge
```

This command purges the system hive and shrinks the size of the registry.

vxverify

`vxverify` is a command that determines if there are any differences between plexes of mirrored volumes. This command compares the volumes on a block by block basis. Blocks that do not match are reported in the output.

Note: For RAID-5 volumes, `vxverify` also checks the integrity of the parity of the volumes.

Typing the following sequence brings up a description of its syntax:

```
vxverify -?
```

```
vxverify [-v] [-p] [-m <f|x>] [-d <n|all>] <DynamicDiskGroupName>  
<Volume Name>[[offset] length]
```

The following attributes apply:

- v Enables verbose output. This specifies that in addition to any blocks that do not match, the output includes the mirror volume name, block number, and number of blocks being compared.
- p Enables the utility to process a partial plex for debugging purposes.
- m Used to specify the mode that should be used to compare the volumes on a block by block basis.

<code>f x</code>	Determines the mode for comparing the mirrored volumes. <ul style="list-style-type: none">■ <code>f</code>- Use to compare all the blocks. Note: This option can be used only when the Smart Move option is not enabled.■ <code>x</code>- Use to compare only the blocks that the file system uses. Note: This option can be used only if the file system is NTFS or ReFS and the mirrors are created using the Smart Move option.
<code>-d</code>	Enables the utility to dump mismatched blocks.
<code>n all</code>	Specifies the number of mismatched blocks to dump. <ul style="list-style-type: none">■ <code>n</code> specifies the number of mismatched blocks.■ <code>"all"</code> specifies the dumping of all mismatched blocks.
<code><DynamicDiskGroupName></code>	Specifies the dynamic disk group of the volume. Note: The name of the disk group is case-sensitive.
<code><Volume Name></code>	Specifies the volume to process.
<code>offset</code>	Specifies the offset for comparison. Offset is specified as a number of blocks.
<code>length</code>	Specifies the length for comparison. Length is specified as a number of blocks. Note: If no length is specified, then the entire volume is compared.

Example

```
vxverify -g Accounting test
```

This command performs a comparison of the volume test of the disk group Accounting and its mirrors.

vxprint

`vxprint` is a command that displays information about SFW and Volume Replicator objects.

Specifying the name of an object results in a display of information for the object. Hierarchies within record associations are displayed in an orderly fashion so that the structure of records is clear. Dashes (-) are displayed in the output wherever there is no available output value. If no option is specified in the command, the command defaults to the -h option.

The typical output format consists of single-line records, each of which include information such as record type, name, object association, object state, length, and other fields. A header line is also written before the record information.

When no disk group is specified with the command, objects in all the disk group are displayed.

Typing the following sequence brings up a description of its syntax:

```
vxprint -?
```

```
vxprint [-h] [-n] [-q] [-l] [-P] [-V] [-v] [-p] [-d] [-s] [-A] [-Q]
[-G] [-E] [-g <DynamicDiskGroupName> ] [Name]
```

The following attributes apply:

-h	List record hierarchies.
-n	List record names.
-q	Suppress the field header in output display.
-l	List record information in verbose format.
-P	List Volume Replicator RLINK records.
-V	List Volume Replicator RVG records.
-v	List volume records.
-p	List plex records.
-d	List disk records.
-s	List subdisk records.
-A	Display all disk groups.
-Q	Suppress the diskgroup header.
-G	List disk groups.
-E	List enclosures.
<-gDynamicDiskGroupName>	Specifies the dynamic disk group.

<Name> Specifies an object to process.

Example

```
vxprint -gAccounting test
```

This command displays information for the volume test in the disk group Accounting.

```
vxprint -v
```

This command displays information of all the volumes on the system.

```
vxprint -p
```

This command displays information of all the plexes on the system.

sfcache

The `sfcache` utility has commands to administer SmartIO.

Table A-1 `sfcache` commands

Command	Description
<p>Creating cache</p> <pre>sfcache create <cachearea_name> [-a] [cache_line_size = <size>] <diskname(s) p#c#t#l#></pre>	<ul style="list-style-type: none"> ■ <i>cachearea_name</i>: Enter a valid name for cache area. ■ -a: Make the cache area auto cache area. ■ <i>diskname(s)</i>: Space separated list of SSDs to be used to create the cache area. ■ p#c#t#l#: Another way to indicate a disk name. The p#c#t#l#parameter corresponds to the port, channel, target, and LUN of a disk.
<p>Deleting cache</p> <pre>sfcache delete <cachearea_name></pre>	<ul style="list-style-type: none"> ■ <i>cachearea_name</i>: Name of the cache area to be deleted.

Table A-1 sfcache commands (continued)

Command	Description
<p>Enabling cache</p> <pre>sfcache -g <DynamicDiskGroupName> enable <VolumeName (s) DriveLetter VmName DrivePath> [-c<cachearea_name>]</pre>	<ul style="list-style-type: none"> ■ -g <DynamicDiskGroupName>: Name of disk group containing source volume. Required when specifying volume name. ■ VolumeName (s): Name of the volume for which caching is to be enabled. ■ DriveLetter: : Drive letter of the volume. ■ VmName: Internal name of the volume; requires the use of the -g option, such as -g DG1 Volume1. ■ DrivePath: A volume that is mounted on an NTFS folder; it is indicated by a path to that folder. ■ -c <cachearea_name>: Name of the cache area to be linked with the volume. If this option is not specified, the volume is linked with the auto cache area.
<p>Disables caching for a specific data volume</p>	<ul style="list-style-type: none"> ■
<p>Situation 1 Brings a cache area online</p>	<p>Brings a cache area online</p>
<p>Situation 2 Takes a cache area offline</p>	<p>Takes a cache area offline</p>
<p>Situation 3 Extends the cache area size</p>	<p>Extends the cache area size</p>
<p>Reduces the cache area size</p>	<p>Reduces the cache area size</p>
<p>Sets the cache area attributes to "auto" or "noauto"</p>	<p>Sets the cache area attributes to "auto" or "noauto"</p>

Table A-1 sfcache commands (*continued*)

Command	Description
Situation 1 Displays the volumes that are associated with a cache area	Displays the volumes that are associated with a cache area
Situation 2 Displays the cache statistics (includes cache hit rate, misses, and, average read and write latencies)	Displays the cache statistics (includes cache hit rate, misses, and, average read and write latencies)

sfcache create

```
sfcache create <cachearea_name> [-a] [cacheline_size = <size>]
<diskname(s) | p#c#t#l#>
```

<i>cachearea_name</i>	Enter a valid name for cache area.
-a	Make the cache area auto cache area.
<i>diskname(s)</i>	Space separated list of SSDs to be used to create the cache area.
<i>p#c#t#l#</i>	Another way to indicate a disk name. The <i>p#c#t#l#</i> parameter corresponds to the port, channel, target, and LUN of a disk.
[<i>cacheline_size</i> = <size>]	Cacheline size or region size in KB. Supported values are 16, 32, and 64. Default value is 64

sfcache delete

```
sfcache delete <cachearea_name>
```

<i>cachearea_name</i>	Name of the cache area to be deleted.
-----------------------	---------------------------------------

sfcache enable

```
sfcache -g <DynamicDiskGroupName> enable
<VolumeName(s) | DriveLetter | VmName | DrivePath> [-c<cachearea_name>]
```

<code>-g</code> <code><DynamicDiskGroupName></code>	Name of disk group containing the source volume. Required when specifying the volume name.
<code>VolumeName (s)</code>	Name of the volume for which caching is to be enabled.
<code>DriveLetter</code>	The drive letter of the volume.
<code>VmName</code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-g DG1 Volume1</code> .
<code>DrivePath</code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<code>-c <cachearea_name></code>	Name of the cache area to be linked with the volume. If this option is not specified, the volume is linked with the auto cache area.

sfcache disable

```
sfcache -g <DynamicDiskGroupName> disable  
<VolumeName (s) | DriveLetter | VmName | DrivePath>
```

<code>-g</code> <code><DynamicDiskGroupName></code>	Name of disk group containing the source volume. Required when specifying the volume name.
<code>VolumeName (s)</code>	Name of the volume for which caching is to be disabled.
<code>DriveLetter</code>	The drive letter of the volume.
<code>VmName</code>	Internal name of the volume; requires the use of the <code>-g</code> option, such as <code>-g DG1 Volume1</code> .
<code>DrivePath</code>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.

sfcache online

```
sfcache online <cachearea_name>
```

`cachearea_name` Name of the cache area to be brought online.

sfcache offline

```
sfcache offline <cachearea_name>
```

cachearea_name Name of the cache area to be brought offline.

sfcache growby

```
sfcache growby <cachearea_name> <diskname(s) | p#c#t#l#>
```

cachearea_name Name of the cache area to be grown.

diskname(s) Space separated list of SSDs to be added to the cache area.

p#c#t#l# Another way to indicate a disk name. The *p#c#t#l#* parameter corresponds to the port, channel, target, and LUN of a disk.

sfcache shrinkby

```
sfcache shrinkby <cachearea_name> <diskname(s) | p#c#t#l#>
```

cachearea_name Name of the cache area to be shrunk.

diskname(s) Space separated list of SSDs to be removed from the cache area.

p#c#t#l# Another way to indicate a disk name. The *p#c#t#l#* parameter corresponds to the port, channel, target, and LUN of a disk.

sfcache set

```
sfcache set <cachearea_name> {-a|-r}
```

cachearea_name Name of the cache area.

-a Set the cache area as auto cache area.

-r Set the cache area as noauto cache area.

sfcache list

```
sfcache list {<[-g <DynamicDiskGroupName>]  
<VolumeName(s) | DriveLetter | VmName | DrivePath> > | <cachearea_name>} [-1]
```

-g<DynamicDiskGroupName> Name of disk group containing the source volume. Required when specifying the volume name.

<i>VolumeName (s)</i>	Name of the volume for which caching is to be disabled.
<i>DriveLetter</i>	The drive letter of the volume.
<i>VmName</i>	Internal name of the volume; requires the use of the -g option, such as -g DG1 Volume1.
<i>DrivePath</i>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<i>cachearea_name</i>	Name of the cache area.
-l	List information in verbose format.

sfcache stat

```
sfcache stat {<[-g <DynamicDiskGroupName>]
<VolumeName (s) | DriveLetter | VmName | DrivePath> > | <cachearea_name>}
{[-l] | [-i<interval>[-n<Count>]]}
```

-g <DynamicDiskGroupName>	Name of disk group containing the source volume. Required when specifying the volume name.
<i>VolumeName (s)</i>	Name of the volume for which caching is to be disabled.
<i>DriveLetter</i>	The drive letter of the volume.
<i>VmName</i>	Internal name of the volume; requires the use of the -g option, such as -g DG1 Volume1.
<i>DrivePath</i>	A volume that is mounted on an NTFS folder; it is indicated by a path to that folder.
<i>cachearea_name</i>	Name of the cache area.
-l	List information in verbose format.
-i <interval>	Displays the statistics every <Interval> seconds. The default and minimum interval is 5 seconds between the instances that the statistics are displayed.
-n<Count>	Stops displaying the interval statistics after <i>Count</i> times.

Tuning SFW

SFW provides the `vxtune` command to tune SFW tunables.

This command is especially useful if you want to experiment with different values to arrive at an optimum value that suits your requirements.

Syntax for `vxtune` command:

```
vxtune [ <tunable> [<value>] ]
```

Parameters for `vxtune` command:

Parameters	Description
------------	-------------

Tunable	Specifies the tunable name whose value you want to display or change.
---------	---

Value	Specifies the value that you want to set for the tunable.
-------	---

A new tunable, `vol_maxkiocount` is applicable for `vxio`. You can tune it as follows:

`vol_maxkiocount`: Specifies the maximum number of parallel I/O operations that SFW can perform.

The maximum permissible value can be 2147483647. The default is 2048.

Displaying the tunable values

Use the `vxtune` command without any parameters to display the value that is currently assigned to the SFW tunable.

Syntax for `vxtune` command:

```
vxtune [<tunable>]
```

To display the default values for the tunable use the command:

```
vxtune
```

To display the value for a specific tunable use the command:

```
vxtune <tunable>
```

The output for `vxtune` command resembles the following:

```
vol_maxkiocount=4096
```

Setting the tunable values

Use the `vxtune` tunable command with the value argument to set the tunable to a new value. After you modify the tunable, the new tunable value is updated immediately.

```
vxtune <tunable> <value>
```

To set the value of `vol_maxkiocount` to 4096:

```
vxtune vol_maxkiocount 4096
```


VDID details for arrays

This appendix includes the following topics:

- [Format of VDID strings for disks belonging to various arrays](#)

Format of VDID strings for disks belonging to various arrays

The Veritas Disk ID (VDID) is a string that contains the information that can identify the array that a disk is contained in. The VDID string is shown in the GUI as a part of the display of information for disks. The VDID string is also shown as a result of the `vxdisk diskinfo` and the `vxdisk list` CLI commands.

The following table describes the format of the VDID strings for the disks that belong to various arrays.

Note: For the disks that the `ddlprov` provider discovers, the VDID string format for its array is "VID_PID_CabinetSerialNumber_LUNSerialNumber". For the disks that the `ddlprov` provider does not discover, the VDID string format for its array is "VID_PID_DISKS_LUNSerialNumber". Note that the following is not a complete list of such arrays.

Array	Model/Type	Format of VDID String
Cisco	MDS	SCSI Inquiry Page 83 data
COMPAQ	HSG80/HSG80CCL	SCSI Inquiry Page 83 data
Default	All models	DevicePath@HostName
EMC	Clarion	EMC_DGC_CabinetSerialNumber_LUNSerialNumber

Format of VDID strings for disks belonging to various arrays

EMC	Symmetrix/DMX	EMC_SYMMETRIX_CabinetSerialNumber_LUNSerialNumber
Hitachi	9960, 9970, 9980, 9200, 9500, 7700, USP	HITACHI_PID_CabinetSerialNumber_LUNSerialNumber
HP-EVA	HSV	HP_PID_CabinetSerialNumber_LUNSerialNumber
HP-XP	1024	HP_1024_CabinetSerialNumber_LUNSerialNumber
HP-XP	512	HP512_CabinetSerialNumber_LUNSerialNumber
HP-XP	128	HP128_CabinetSerialNumber_LUNSerialNumber
IBM	DS8000, DS6000	IBM_PID_CabinetSerialNumber_LUNSerialNumber
IBM	200, 700, 600, 400, Fast_T	IBM_PID_CabinetSerialNumber_LUNSerialNumber
IBM-LSI	All models	SCSI Inquiry Page 83 data
IBM Shark	F20, ESS 800, F10	IBM_PID_CabinetSerialNumber_LUNSerialNumber
JBOD	all models	{WWWN}+{SCSI LUNID}
NETAPP	F825, FAS900	NETAPPP_LUN_CabinetSerialNumber_LUNSerialNumber
SUN-T3	All models	SCSI Inquiry Page 83 data
3PAR	Inserv Storage Server, E-Class, F-Class, S-Class, T-Class	3PARdata_PID_CabinetSerialNumber_LUNSerialNumber
DELL	Equallogic PS5000/PS6000 Series	VID_PID_CabinetSerialNumber_LUNSerialNumber
PILLAR	Axiom 300/500/600	VID_PID_CabinetSerialNumber_LUNSerialNumber
Fujitsu	All models	VID_PID_DISKS_LUNSerialNumber
Compellent	All models	VID_PID_DISKS_LUNSerialNumber
EMC	VPLEX	VID_PID_DISKS_LUNSerialNumber
HP	MSA P2000	VID_PID_DISKS_LUNSerialNumber
HP	MSA 2040	VID_PID_DISKS_LUNSerialNumber

IBM	DS5020	VID_PID_DISKS_LUNSerialNumber
IBM	DS3524	VID_PID_DISKS_LUNSerialNumber
Sun	6540	VID_PID_DISKS_LUNSerialNumber
Sun	7310	VID_PID_DISKS_LUNSerialNumber

Glossary

active/active	In Storage Foundation for Windows, the mode in which DMP allocates the data transfer across the possible paths to and from an array, thus enabling load balancing. With this mode, DMP implements a round-robin algorithm, selecting each path in sequence for each successive data transfer to or from a disk. For example, if paths A and B are active, the first disk transfer occurs on path A, the next on path B, and the next on path A again.
active partition or volume	<p>The partition or volume from which the computer starts up. On a basic disk, the active partition must be a primary partition. For a dynamic disk, a dynamic volume cannot be marked active directly. However, a basic disk with a system partition can be upgraded to dynamic. After the disk is upgraded and the computer is restarted, the system partition becomes a system volume, which retains its active status.</p> <p>Upgrading a basic disk with a system partition to a dynamic disk preserves the partition table information, which is required when reinstalling the Windows operating system.</p>
active/passive	In Storage Foundation for Windows, the mode in which DMP allocates data transfer to and from an array across the preferred path. This path is always active, and the other path or paths act as backups that are called into service if the current operating path fails. This option does not provide load balancing.
boot partition volume	The volume that contains the Windows operating system and its support files. The boot volume can be in the same location as the system volume. In Windows, the system partition or volume is the one the computer starts from, while the boot volume is the one that has the operating system and support files.
critical threshold	In Storage Foundation for Windows, the percentage of volume capacity at which an error is reported. The default is 90 percent.
disk	A physical data storage device attached to a computer.
disk evacuation	The process of evacuating or moving the entire contents of a healthy disk to the free space on a dynamic disk in the same dynamic disk group.
disk group	Storage Foundation for Windows organizes disks into disk groups. Disk groups provide a way of organizing disks in a system and simplifying storage management for systems with large numbers of disks. They also allow disks to be moved between computers to easily transfer the storage between computers.

disk replacement	In Storage Foundation for Windows, the process of replacing a failed disk with an empty basic disk. The volume configuration is recreated on the new disk. The contents of nonredundant volumes are not guaranteed. Redundant volumes will be automatically resynchronized.
drive path	The path to a local drive that is mounted at an empty folder on an NTFS volume.
dynamic disk	A dynamic disk is a physical disk that can contain dynamic volumes created with Storage Foundation for Windows. A dynamic volume organizes space on one or more physical disks by using a specific type of volume layout. The six types of dynamic volume layouts are simple, spanned, mirrored, striped, RAID-5, and mirrored striped (RAID 0+1). On a dynamic disk, space is organized through volumes rather than partitions. Because a dynamic disk does not have the partitioning scheme used by Windows NT, Windows 95/98, and MS-DOS, dynamic disks cannot be accessed through those operating systems.
Dynamic Multipathing	In Storage Foundation for Windows, the Dynamic Multipathing option adds fault tolerance to disk storage by making use of multiple paths between a computer and individual disks in an attached disk storage system. Disk transfers that would have failed because of a path failure are automatically rerouted to an alternate path. Dynamic Multipathing also improves performance by allowing load balancing between the multiple paths. Two forms of Dynamic Multipathing are available, DMP and MPIO.
dynamic volume	In Storage Foundation for Windows, dynamic volumes are volumes created on dynamic disks in place of partitions. A dynamic volume consists of a portion or portions of one or more physical disks and is organized in one of five volume layout types: concatenated, mirrored, striped, RAID-5, and mirrored striped (RAID 0+1). The size of a dynamic volume can be increased if the volume is formatted with NTFS and there is unallocated space on a dynamic disk within the dynamic disk group onto which the volume can be extended.
extended partition	A portion of a basic disk that can contain logical drives. An extended partition is used to have more than four volumes on a basic disk. A basic disk can contain up to four primary partitions or three primary partitions plus an extended partition. The extended partition can be further divided into up to 32 logical drives.
FR (FastResync)	In Storage Foundation for Windows, FastResync is a part of the FlashSnap feature. FR supports resynchronization of mirrors by copying only changes to the temporarily split mirror by using FR logging. This reduces the time it takes to rejoin a split mirror to the mirror set and also reduces the server CPU cycles needed to complete the resynchronization. FastResync can also be used on ordinary mirrored volumes to speed up resynchronization.
failover	An operation in which the failure of one appliance, program, or security gateway causes another to pick up its workload automatically.

FT (fault tolerant) disk	Refers to a disk that contains fault tolerant volumes that are created in Windows NT Disk Administrator.
FT (fault tolerant) volume	Refers to a RAID volume that was originally created in Windows NT Disk Administrator. These volumes include spanned volumes (volume sets), mirrored volumes (mirror sets), striped volumes (stripe sets), and RAID-5 volumes (stripe sets with parity). The FT refers to fault tolerant, even though some of the volume sets are not fault tolerant. Windows Server 2003 does not support FT volumes.
hot relocation	In Storage Foundation for Windows, when a disk fails, hot relocation automatically moves all subdisks from redundant volumes on the failed disk to hot spare disks, or to free space on other disks if enough space is not available on hot spare disks.
hot spot	In Storage Foundation for Windows, a hot spot is an area of high I/O activity that may cause bottlenecks in I/O throughput.
logical drive	<p>A logical drive is a simple volume that resides on an extended partition on a basic disk. Logical drives are limited to the space on the extended partition. They cannot span multiple disks. A logical drive can be formatted and assigned a drive letter.</p> <p>An extended partition can be subdivided into as many as 32 logical drives. All or part of the free space in an extended partition can be used when creating logical drives.</p>
LUN (logical unit number)	The number that, when combined with the target ID, uniquely identifies a disk on the port.
monitor interval	In Storage Foundation for Windows, DMP (Dynamic Multipathing) monitors the paths to an array to determine whether they are functioning properly. The monitor interval specifies the time interval for the monitoring of that array.
NTFS (Windows NT file system)	An advanced file system designed for use specifically within Windows operating systems. It supports file system recovery, extremely large storage media, long file and folder names, and file and folder compression.
parity	<p>Redundant information that is associated with a block of information. Parity is a calculated value used to reconstruct data after a failure.</p> <p>RAID-5 volumes stripe data and parity intermittently across a set of disks. Within each stripe, the data on one disk is parity data and the data on the other disks is normal data. RAID-5 volumes, therefore, require at least three disks to allow for this extra parity information. When a disk fails, the parity information on the good disks can be used to recreate the data on the failed disk.</p>
partition	A portion of a physical disk that functions as though it were a physically separate disk. Partitions can be created only on basic disks. Partitions cannot span disks; they must be contiguous regions. When a basic disk that contains partitions is upgraded to a dynamic disk, the partitions become simple volumes on the dynamic disk.

polling interval	In Storage Foundation for Windows, the polling interval is the interval at which volume capacities are checked.
primary disk group	The disk group that contains the computer's boot or system disk. All other dynamic disk groups are called secondary disk groups. If none of the dynamic disk groups on a system contain the boot or system disk, then there is not a primary disk group.
primary partition	A volume that is created by using unallocated space on a basic disk. Microsoft Windows NT and other operating systems can start from a primary partition. Up to four primary partitions can be created on a basic disk, or three primary partitions and an extended partition. Primary partitions can be created only on basic disks and cannot be subpartitioned. However, the extended partition can be further divided into as many as 32 logical drives.
ReFS (Resilient File System)	An advanced file system that is built on the foundation of NTFS. ReFS is supported on Windows Server 2012 and 2012 R2 operating systems only.
secondary disk group	Any dynamic disk group that is not primary. A primary disk group is a dynamic disk group that contains the computer's boot or system disk. All other dynamic disk groups are called secondary disk groups.
simple dynamic volume	<p>A simple dynamic volume consists of a single contiguous region (or subdisk) on a single physical disk. Simple dynamic volumes can be extended or mirrored. A simple volume can be extended within the same disk or onto additional disks.</p> <p>When a basic disk with a partition is upgraded, the partition becomes a simple volume. An extended partition on a basic disk also becomes a simple volume when the disk is upgraded to dynamic.</p>
spanned volume	A volume made up of disk space on more than one physical disk. Space can be added to a spanned volume by extending it at any time. Spanned volumes can be created only on dynamic disks. Spanned volumes by themselves are not fault tolerant. However, they can be mirrored to be made fault tolerant.
statistics threshold	In Storage Foundation for Windows, the statistics threshold is the minimum amount of I/O per second that a disk or subdisk must have to display an I/O status other than the Low I/O activity status.
striped volume (RAID-0)	A volume that stores data in stripes on two or more physical disks. Data in a striped volume is allocated alternately and evenly (in stripes) to the disks of the striped volume. Striped volumes can be created only on dynamic disks. Striped volumes by themselves are not fault tolerant; however, they can be mirrored to be made fault tolerant. They also can be extended.
subdisk	Refers to a region of contiguous space on a disk. In Storage Foundation for Windows, subdisks are the basic units in which disk space in a volume is allocated. A disk can be divided into one or more subdisks. Each subdisk represents a specific portion of the volumes on a dynamic disk.

A Storage Foundation for Windows disk may contain multiple subdisks, but subdisks cannot overlap or share the same portions of a disk. Any disk space that is not part of a subdisk is considered to be unallocated space, which can be used to create new volumes.

Subdisks can be moved to improve disk performance.

system partition or volume	The partition or volume that has the files needed to load the operating system. It is the same as the active partition or volume. In Windows, the system partition or volume is the one the computer starts from, while the boot volume is the one that has the operating system and support files. The system partition or volume and the boot partition or volume can be the same partition or volume.
target ID	The number that, when combined with the LUN, uniquely identifies a disk on the port.
volume	A volume is a logical entity that is made up of a portion or portions of one or more physical disks. A volume can be formatted with a file system and can be accessed by a drive letter or a mount path. Like disks, volumes can be basic or dynamic.
volume read policy	The volume read policy on a dynamic volume allows an administrator to specify either that a particular mirror be used for reads or that all mirrors be read in turn in round-robin fashion for each nonsequential I/O detected.
warning threshold	In Storage Foundation for Windows, the percentage of volume capacity at which a warning message is sent. The default is 80 percent of volume capacity.