

Veritas Storage Foundation™ and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange

Windows Server 2003
Windows Server 2008

5.1 Service Pack 2



Veritas Storage Foundation and HA Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange

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

Product version: 5.1. Service Pack 2

Document version: 5.1.SP2.0

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

This Symantec product may contain third party software for which Symantec 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. Please see the Third Party Legal Notice file accompanying this Symantec product for more information on the Third Party Programs.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information

- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

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

North America and Latin America supportsolutions@symantec.com

Contents

Section 1 Introduction

Chapter 1 Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange

About the solutions guides	20
Quick Recovery solutions	20
Microsoft clustering solutions	20

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	21
Starting the Configuration Center	22
Available options from the Configuration Center	23
About running the Configuration Center wizards	25
Following the workflow in the Configuration Center	27
Solutions wizard logs	28

Section 2 Quick Recovery

Chapter 3 Quick Recovery concepts

About Quick Recovery snapshot solutions	31
Advantages of Quick Recovery snapshots	32
Quick Recovery process	33
About the components used in Quick Recovery	33
FlashSnap and FastResync	33
Integration with Microsoft Volume Shadow Copy Service	34

Chapter 4 Planning a Quick Recovery snapshot solution

System requirements	37
Supported software	38
Storage requirements	39

Methods of implementing Quick Recovery snapshots	41
About the Quick Recovery Configuration Wizard	41
About the VSS Exchange Snapshot Scheduler Wizard	42
About the VSS Snapshot and Snapback wizards and the vxsnap utility	43
Planning your Quick Recovery solution	44
Backup types for snapshot sets	44
About logs	44
Recommendations and best practices	45
Recommendations for Exchange storage configuration	46
Recommendations for maintenance and backups	46
VCS cluster considerations	47
Microsoft cluster considerations	48
VVR considerations	50
Exchange 2007 passive copy snapshot considerations	50

Chapter 5 Configuring Exchange for Quick Recovery snapshots

Tasks for configuring Exchange storage	54
Reviewing the prerequisites	54
Reviewing the configuration	55
Creating dynamic disk groups	56
Creating dynamic volumes	57
Pointing the databases and log paths to the SFW volumes (Exchange 2003)	60
Moving the First Storage Group (Exchange 2003)	62
Pointing the databases and log paths to the SFW volumes (Exchange 2007)	65
Moving the First Storage Group (Exchange 2007)	67

Chapter 6 Implementing snapshot sets with the configuration wizard

About the Quick Recovery Configuration Wizard	69
About snapshot sets	70
About snapshot templates	71
Tasks for implementing snapshot sets with the configuration wizard	73
Reviewing the prerequisites	73
Scheduling and creating snapshot sets	74
Viewing the status of scheduled snapshots	83
Troubleshooting scheduled snapshots	84
Deleting or modifying schedules	86
Synchronizing schedules after adding a cluster node	87

Chapter 7 Scheduling or creating an individual snapshot set

- About scheduling or creating an individual snapshot set 89
- Tasks to schedule a new snapshot 90
- Tasks to create a one-time snapshot set 91
- Reviewing the prerequisites 91
- Preparing the snapshot mirrors 92
- Scheduling a snapshot set 93
 - Viewing the status of scheduled snapshots 100
- Troubleshooting scheduled snapshots 101
 - Deleting schedules 103
- Creating a one-time snapshot set 104
- Refreshing a snapshot set manually 107
 - Reattaching the split-mirror snapshots 108
 - Creating the refreshed snapshot set 110

Chapter 8 Recovering Exchange mailbox storage group or databases

- About recovery using Quick Recovery snapshots 111
- Tasks for recovery using Quick Recovery snapshots 112
- Prerequisites for recovery 113
- Recovery using an Exchange 2007 passive copy snapshot 113
 - Overview of recovery steps using an Exchange 2007 passive copy snapshot 114
 - Example of restoring from a passive copy 115
- Recovering using the VSS Restore Wizard 115
 - Recovering using the vxsnap utility 120
- Post-recovery steps 123

Chapter 9 Recovering after hardware failure

- About recovery after hardware failure 125
 - Recovery scenarios 126
 - Recovery tasks 126
- Tasks for recovering after hardware failure 127
- Reviewing the prerequisites 128
- Reviewing the sample configuration 128
- Scenario I: Database and transaction logs volumes are missing 129
 - Identifying the missing volumes 131
 - Dismounting the Exchange database 131
 - Deleting missing volumes from Storage Foundation for Windows ... 133
 - Replacing hardware and adding disks to the dynamic disk group 133
 - Changing the drive letter or mount points of the snapshot volumes 134

- Restoring the storage group to the point in time 136
- Refreshing the snapshot set 137
- Scenario II: Database volumes missing, transaction logs are available 137
 - Identifying the missing volumes 138
 - Dismounting Exchange databases 139
 - Deleting missing volumes from Storage Foundation for Windows ... 139
 - Replacing hardware and adding disks to the dynamic disk group 140
 - Changing the drive letter or mount points of the
 - snapshot volumes 141
 - Restoring the storage group to the Point in Time (PIT) 143
 - Recovering the storage group to the Point of Failure (POF) 144
 - Refreshing the snapshot set 145
- Scenario III: Some DB volumes missing, transaction logs are available ... 145
 - Identifying the missing volumes 146
 - Dismounting Exchange databases 147
 - Deleting missing volumes from Storage Foundation for Windows ... 148
 - Replacing hardware and adding disks to the dynamic disk group 148
 - Changing the drive letter or mount points of the
 - snapshot volumes 149
 - Recovering the storage group to the Point of Failure (POF) 151
 - Refreshing the snapshot set 152
- Refreshing the snapshot set 153
- Refreshing the snapshot set on the current disks 153
 - Reattaching healthy snapshot volumes 154
 - Clearing the snapshot association for volumes whose
 - drive letters or mount points were reassigned 155
 - Creating snapshot mirrors of volumes whose
 - drive letters or mount points were reassigned 155
 - Creating the new snapshot set 156
- Moving the production volumes to different disks
 - and refreshing the snapshot set 160
 - Reattaching healthy snapshot volumes 160
 - Clearing the snapshot association for volumes whose
 - drive letters or mount points were reassigned 161
 - Adding mirrors to volumes whose drive letters or
 - mount points were reassigned 162
 - Creating snapshot mirrors of volumes whose drive letters
 - or mount points were reassigned 162
 - Creating the new snapshot set 163

Chapter 10	Vxsnap utility command line reference	
	About the vxsnap utility	169
	Vxsnap keywords	169
	vxsnap prepare	171
	vxsnap create	172
	vxsnap reattach	175
	vxsnap restore	176
Section 3	Microsoft Clustering Solutions	
	About Microsoft clustering solutions	181
Chapter 11	Deploying SFW with MSCS: New Exchange installation	
	Tasks for a new Exchange installation with SFW and MSCS (Windows Server 2003)	184
	Reviewing the requirements	186
	Supported software	186
	Disk space requirements	187
	System requirements	187
	Reviewing the configuration	188
	Configuring the storage hardware and network	190
	Establishing an MSCS cluster	191
	Creating the MSDTC resource (Exchange 2003)	192
	Installing SFW with MSCS/Failover Cluster option	194
	SFW installation tasks	194
	Pre-installation tasks	194
	Installing Veritas Storage Foundation for Windows	196
	Post-installation tasks	199
	Configuring SFW disk groups and volumes	201
	Planning disk groups and volumes	201
	Creating a disk group	202
	Creating volumes	204
	Preparing the forest and domain (Exchange 2003)	208
	Adding a Volume Manager Disk Group resource for Exchange 2007 installation	208
	Installing Exchange Server	209
	Creating an Exchange virtual server group (Exchange 2003)	210
	Adding Volume Manager Disk Group resources to the Exchange 2007 group	220
	Moving Exchange databases and logs to shared storage (Exchange 2003)	221

Moving Exchange databases and logs to shared storage (Exchange 2007)	223
Implementing a dynamic mirrored quorum resource	225
Creating a dynamic cluster disk group for the quorum resource with mirrored volume	226
Creating the quorum resource for the cluster group	226
Changing the quorum resource to a dynamic mirrored quorum resource	228
Verifying the cluster configuration	228

Chapter 12 Deploying SFW with Microsoft failover clustering: New Exchange installation

Tasks for a new Exchange installation with SFW and failover clustering (Windows Server 2008)	232
Reviewing the requirements	233
Supported software for Microsoft failover clusters with SFW	234
Disk space requirements	235
System requirements	236
Reviewing the configuration	236
Configuring the storage hardware and network	238
Establishing a Microsoft failover cluster	239
Installing SFW with MSCS/Failover Cluster option	242
SFW installation tasks	242
Pre-installation tasks	242
Installing Veritas Storage Foundation for Windows	243
Post-installation tasks	246
Configuring SFW disk groups and volumes	248
Planning disk groups and volumes	248
Creating dynamic cluster disk groups	249
Creating volumes	251
Managing disk groups and volumes	255
Importing a disk group and mounting a volume	255
Unmounting a volume and deporting a disk group	256
Implementing a dynamic mirrored quorum resource	256
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	257
Adding the Volume Manager Disk Group resource for the quorum ..	257
Changing the quorum resource to a dynamic mirrored quorum resource	258
Adding a Volume Manager Disk Group resource for Exchange 2007 installation	259
Installing Exchange Server	259

- Adding the Volume Manager Disk Group resources to the Exchange group260
- Setting the database dependency on the disk group resource260
- Moving Exchange databases and logs to shared storage261
- Verifying the cluster configuration263

Chapter 13

Deploying SFW with MSCS and Exchange in a campus cluster

- Tasks for a new Exchange, SFW and MSCS installation in a campus cluster (Windows Server 2003)266
- Reviewing the requirements269
 - Supported software269
 - System requirements270
 - Disk space requirements271
- Reviewing the configuration272
 - Overview of campus clustering with MSCS273
 - MSCS campus cluster failure scenarios275
 - MSCS quorum and quorum arbitration278
- Configuring the network and storage280
- Establishing an MSCS cluster282
 - Installing and configuring the operating system and MSCS on Server A282
 - Configuring the shared storage and creating a partition for the Cluster quorum disk283
 - Creating the first node of the cluster on Server A283
 - Installing and configuring the operating system and MSCS on Server B283
 - Connecting the two nodes283
 - Creating the second node of the cluster on Server B284
 - Verifying the cluster configuration284
- Creating the MSDTC resource (Exchange 2003 only)285
- Installing SFW286
 - SFW installation tasks286
 - Pre-installation tasks287
 - Installing Veritas Storage Foundation for Windows289
 - Post-installation tasks292
- Creating disk groups and volumes294
 - Configuring the disks and volumes295
 - Creating a dynamic (cluster) disk group296
 - Creating a volume298

Implementing a dynamic quorum resource	303
Creating a dynamic cluster disk group for the quorum, mirrored	303
Making the quorum cluster disk group an MSCS resource	304
Changing the quorum resource to the dynamic mirrored quorum resource	305
Adding a Volume Manager Disk Group resource for Exchange 2007 installation	306
Installing the application on the cluster nodes	307
Setting up a group for Exchange 2003 in MSCS	309
Adding Volume Manager Disk Group resources to the Exchange 2007 group	310
Moving Exchange databases and logs to dynamic volumes	310
Moving Exchange 2003 databases and logs	310
Moving Exchange 2007 databases and logs	313
Verifying the cluster configuration	314

Chapter 14 Deploying SFW with Microsoft failover clustering and Exchange in a campus cluster

Tasks for a new Exchange, SFW and Microsoft failover clustering installation in a campus cluster (Windows Server 2008)	318
Reviewing the requirements	320
Supported software	320
System requirements	321
Disk space requirements	322
Reviewing the configuration	323
Overview of campus clustering with Microsoft clustering	325
Campus cluster failure with Microsoft clustering scenarios	326
Microsoft clustering quorum and quorum arbitration	330
Configuring the network and storage	331
Establishing a Microsoft failover cluster	334
Connecting the two nodes	335
Installing SFW	336
SFW installation tasks	336
Pre-installation tasks	336
Installing Veritas Storage Foundation for Windows	337
Post-installation tasks	340
Creating disk groups and volumes	342
Configuring the disks and volumes	343
Creating a dynamic (cluster) disk group	344
Creating a volume	346

- Implementing a dynamic quorum resource351
 - Creating a dynamic cluster disk group and a mirrored volume
 - for the quorum resource351
 - Adding the Volume Manager Disk Group resource for the quorum ..352
 - Changing the quorum resource to the dynamic mirrored
 - quorum resource353
 - Adding a Volume Manager Disk Group resource
 - for Exchange 2007 installation354
- Installing the application on the cluster nodes355
- Adding the Volume Manager Disk Group resources
 - to the Exchange group356
- Setting the database dependency on the disk group resource356
- Moving the Exchange databases and logs to the dynamic volumes357
- Verifying the cluster configuration359

Chapter 15 Deploying SFW and VVR with MSCS: New Exchange installation

- Tasks in a new Exchange installation with SFW,
 - VVR, and MSCS (Windows Server 2003)362
- Reviewing the requirements365
 - Supported software365
 - Disk space requirements367
 - System requirements367
- Reviewing the configuration368
- Configuring the primary site371
- Installing SFW with MSCS/Failover Cluster option371
 - SFW installation tasks372
 - Pre-installation tasks372
 - Installing Veritas Storage Foundation for Windows374
 - Post-installation tasks377
- Completing the primary site configuration381
- Setting up the secondary site (Exchange 2003)382
- Setting up the secondary site (Exchange 2007)382
- Installing Exchange on the secondary site (Exchange 2007)384
- Setting up the Exchange group on the secondary site (Exchange 2007) ..385
- Moving the Volume Manager Disk Group resources
 - to the Exchange group (Exchange 2007)386
- VVR components overview387
- Creating resources for VVR (primary and secondary sites)387
- Setting up the replicated data sets (RDS) for VVR389
- Creating the RVG resource (primary and secondary sites)400
 - Setting the System Attendant resource dependency
 - on the RVG resource (Exchange 2003)401

Setting the database resource dependency	
on the RVG resource (Exchange 2007)	402
Normal operations and recovery procedures	404
Normal operations	404
Performing planned migration	404
Replication recovery procedures	405

Chapter 16 Deploying SFW and VVR with Microsoft failover clustering: New Exchange installation

Tasks for a new Exchange, SFW, VVR, and failover clustering installation (Windows Server 2008)	410
Reviewing the requirements	412
Supported software for Microsoft failover clusters with SFW	412
Disk space requirements	413
System requirements	415
Reviewing the configuration	416
Configuring the primary site	419
Installing SFW with MSCS/Failover Cluster option	419
SFW installation tasks	419
Pre-installation tasks	420
Installing Veritas Storage Foundation for Windows	420
Post-installation tasks	424
Completing the primary site configuration	429
Setting up the secondary site	430
Installing Exchange on the secondary site	431
Setting up the Exchange group on the secondary site	432
Moving the Volume Manager Disk Group resources	
to the Exchange group	433
VVR components overview	434
Creating resources for VVR (primary and secondary sites)	434
Setting up the replicated data sets (RDS) for VVR	435
Creating the RVG resource (primary and secondary sites)	446
Setting the database resource dependency	
on the RVG resource	447
Normal operations and recovery procedures	449
Normal operations	449
Performing planned migration	449
Replication recovery procedures	450

Introduction

- [Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange](#)
- [Using the Solutions Configuration Center](#)

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft Exchange

This chapter includes the following topics

- [About the solutions guides](#)
- [Quick Recovery solutions](#)
- [Microsoft clustering solutions](#)

About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange* contains Exchange 2003 and 2007 solutions using the following:

- [Quick Recovery solutions](#)
- [Microsoft clustering solutions](#)

Solutions for Microsoft Exchange 2003 and 2007 for high availability and disaster recovery are in *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange*.

Separate guides are available for Exchange 2010, for Microsoft SQL solutions and for other application solutions.

Quick Recovery solutions

Quick Recovery is the process of creating and maintaining on-host point-in-time images of production Exchange storage groups that can be used to quickly recover missing or corrupt Exchange storage groups or individual database stores.

These images, called snapshot sets, contain snapshots of the volumes associated with the databases and transaction logs of an Exchange storage group. Quick Recovery is designed to augment your traditional backup methodology.

Microsoft clustering solutions

Microsoft clustering solutions are covered in separate chapters according to operating system:

- Microsoft Cluster Server (MSCS) on Windows Server 2003
- Microsoft failover clustering on Windows Server 2008

Microsoft clustering solutions may be used with Veritas Storage Foundation for Windows to provide high availability for Exchange.

Veritas Storage Foundation for Windows and Veritas Volume Replicator may be used with Microsoft clustering to provide replication support for Exchange.

Quick Recovery

- Quick Recovery concepts
- Planning a Quick Recovery snapshot solution
- Configuring Exchange for Quick Recovery snapshots
- Implementing snapshot sets with the configuration wizard
- Scheduling or creating an individual snapshot set
- Recovering Exchange mailbox storage group or databases
- Recovering after hardware failure
- Vxsnap utility command line reference

Quick Recovery concepts

This chapter includes the following topics:

- [About Quick Recovery snapshot solutions](#)
- [Advantages of Quick Recovery snapshots](#)
- [Quick Recovery process](#)
- [About the components used in Quick Recovery](#)

About Quick Recovery snapshot solutions

Veritas Storage Foundation for Windows (SFW) provides the capability to create point-in-time images of production Exchange storage groups. These images, called snapshot sets, contain snapshots of the volumes associated with the databases and transaction logs of an Exchange storage group. A snapshot set is a complete copy of the storage group at the specific point in time the snapshots are taken.

SFW uses Veritas FlashSnap™ technology along with the Microsoft Volume Shadow Copy Service (VSS) framework to quiesce an Exchange storage group and ensure a persistent snapshot of the production data.

Quick Recovery is the term for creating and maintaining the snapshot sets on-host for use in quickly recovering missing or corrupt Exchange storage groups or individual database stores. The Quick Recovery solution provides fast recovery from logical errors and eliminates the time-consuming process of restoring data from tape. Storage groups can be recovered to the point in time when the snapshot was taken or, by using current logs, rolled forward to the point of failure.

Using the SFW Quick Recovery Configuration wizard, you can create multiple snapshot sets for each storage group and set up schedules for creating and refreshing the snapshot sets. The snapshot sets can be maintained on-host as a Quick Recovery solution.

If you are using Veritas Volume Replicator (VVR) for replication, you can also synchronize snapshot sets on the secondary site. See *Veritas Volume Replicator, Administrator's Guide*.

SFW snapshot solutions use a split-mirror snapshot method. The snapshot is a separate persistent volume that contains an exact duplicate of all the data on the original volume at the time the snapshot is taken. This type of persistent physical snapshot is also known as a Clone (HP) or a BCV (EMC). In contrast, copy-on-write snapshots, also known as metadata snapshots, only copy changed blocks to the snapshot and do not create a separate physical volume.

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the storage group, the snapshot set requires the same amount of space as the original volumes.

Veritas FlashSnap technology is also integrated into the Veritas NetBackup 6.0 Advanced Client Option and Symantec Backup Exec 10d Advanced Disk-based Backup Option. These products are the preferred solution for on and off host snapshot-assisted backup.

Advantages of Quick Recovery snapshots

A Quick Recovery solution serves as a first line of defense to recover corrupted or missing Exchange storage groups or databases. Maintaining a snapshot set requires just the few seconds it takes to detach a split-mirror snapshot from its original volume. On-host snapshot recovery is faster than restoring a full backup from tape or other media; on-host snapshot recovery reduces downtime and helps meet service-level agreements for application availability.

In addition to the primary benefit of recovery from logical errors, snapshot sets can be moved over a SAN to another server and used for other purposes including:

- Application tuning and testing—data can be updated and modified in a realistic environment without impacting users.
- Business reporting and decision analysis—up-to-date data is available with minimal impact on the production environment.

Quick Recovery process

The Quick Recovery process can be broken down into the following phases:

- Creating an initial snapshot set
This has two stages:
 - Preparing the mirror for the snapshot set
This stage takes a while and should be scheduled for a time of low activity.
 - Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again, as needed or according to a pre-set schedule
This stage is automated by setting up snapshot schedules using the Quick Recovery wizard or VSS Exchange Snapshot Scheduler wizard.
- Using a snapshot set to recover a storage group or corrupted database

About the components used in Quick Recovery

SFW snapshot solutions use Veritas FlashSnap and FastResync technology along with the Microsoft Volume Shadow Copy Service framework.

FlashSnap and FastResync

Veritas FlashSnap provides the ability to create and maintain the on-host point-in-time copies of volumes that are integral to the snapshot solutions. Both the original and snapshot volume may consist of multiple physical devices, as in the case of RAID 0+1 (Mirrored Striped) volumes. FlashSnap cannot be used with software RAID-5 volumes.

FastResync is a FlashSnap feature that optimizes the resynchronization of a snapshot volume and its original volume. FlashSnap uses FastResync technology to track the changed blocks in an original volume after a snapshot is detached. A Disk Change Object (DCO) volume is automatically created to store a record of these changes. When the snapshot volume is resynchronized with the original volume, only the changed data blocks are written to the snapshot volume. This greatly reduces the time and performance impact of resynchronization which means that a snapshot set can be refreshed with minimal impact to production.

Integration with Microsoft Volume Shadow Copy Service

SFW integrates with the Microsoft Volume Shadow Copy Service (VSS) as both a VSS Requestor and a VSS Provider. This integration is provided by FlashSnap.

The Volume Shadow Copy Service (VSS) process allows the databases of an Exchange Server storage group to be frozen before the snapshot operation occurs and then thawed immediately after it. This quiescing, supported by Exchange Server 2003 and higher at the storage group level, allows for Microsoft supported and guaranteed persistent snapshots of your data.

FlashSnap integrates with VSS to create a snapshot set containing snapshot volumes of all the volumes associated with an Exchange storage group without taking the databases offline.

The VSS framework

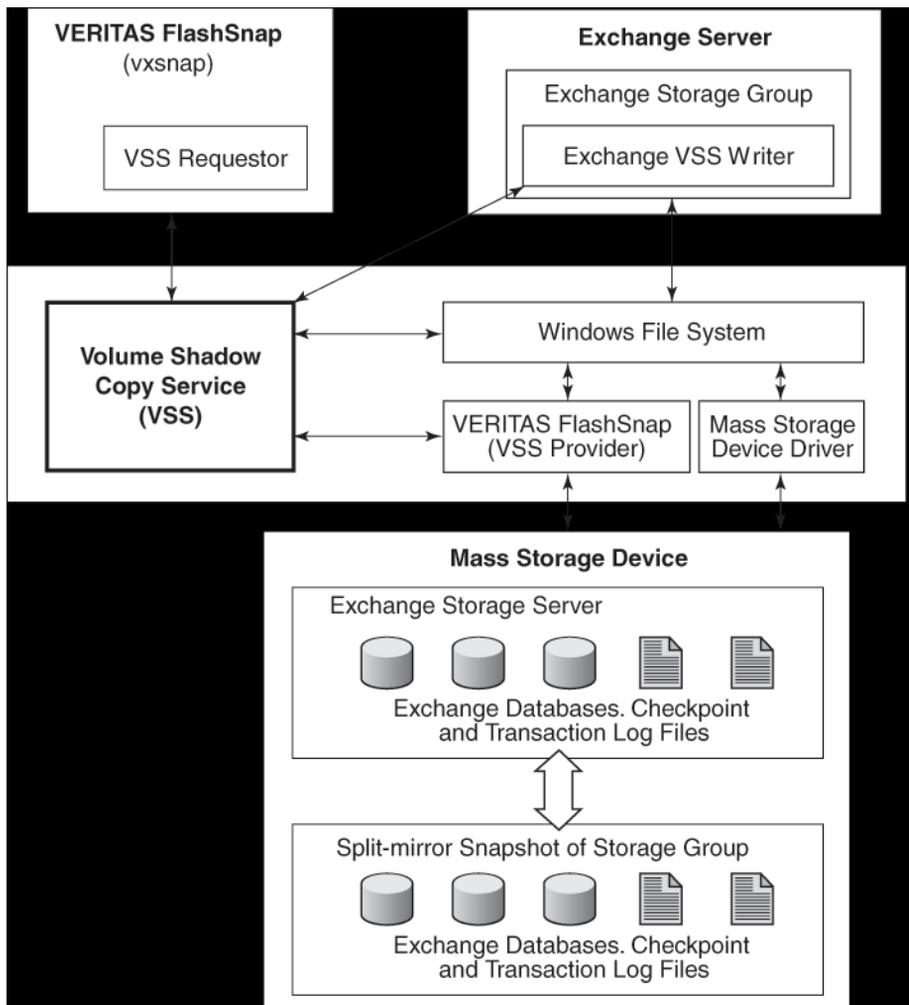
There are four components to the VSS framework: Requestor, Writer, Provider, and the Volume Shadow Copy Service itself.

Table 3-1 VSS framework components

Component	Action
Volume Shadow Copy Service	Talks to and coordinates the Requestor, Provider, and Writer.
Requestor	As a Requestor, the vxsnap component of FlashSnap notifies the VSS coordinator to initiate the VSS request to prepare Exchange for quiescing and later requests that the snapshot process begin.
Writer	As Writers, VSS-enabled applications such as Exchange respond to requests to prepare and participate in the generation of snapshots, provide success/failure status, and provide information about the application including what is to be backed up and restored, and restore strategy. The Exchange 2007 VSS Writer has two instances, the Store Writer and the Replica Writer. The Replica Writer is available for use if either Local Continuous Replication (LCR) or Cluster Continuous Replication (CCR) is enabled. The Replica Writer is supported by SFW Flashsnap for snapshot operations but not for scheduling or restore operations.
Provider	As a Provider, FlashSnap creates the persistent snapshot.

Figure 3-1 shows the steps in the Volume Shadow Copy Service Process.

Figure 3-1 The Volume Shadow Copy Service process



FlashSnap integrates with Volume Shadow Copy Service as both a VSS Requestor and a VSS Provider.

The following steps occur, in the order presented, 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.
- The VSS coordinator service calls the Exchange 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 to quiesce and freeze the databases in preparation for the snapshot. When this is accomplished, the Exchange 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. The snapshot process takes a maximum of 10 seconds. After the snapshot volume is detached, the FlashSnap Provider informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the Exchange Writer to thaw (release the freeze) and resume normal I/O.

The following steps occur, in the order presented, during the restore process:

- Acting as a VSS Requestor, FlashSnap notifies the VSS coordinator service to prepare for a restore operation.
- The VSS coordinator service calls the Exchange VSS Writer, which prepares for the restore operation.
- The FlashSnap utility restores the snapshot volumes. After the snapback operation completes, the FlashSnap utility informs the VSS coordinator service to proceed.
- The VSS coordinator service notifies the Exchange Writer to process the post-restore operations.

Planning a Quick Recovery snapshot solution

This chapter includes the following topics:

- [System requirements](#)
- [Methods of implementing Quick Recovery snapshots](#)
- [Planning your Quick Recovery solution](#)
- [Backup types for snapshot sets](#)
- [About logs](#)
- [Recommendations and best practices](#)

System requirements

A Veritas Storage Foundation for Windows (SFW) Quick Recovery solution can be implemented on either a standalone system or a clustered system. Quick Recovery snapshots are supported in both Veritas Cluster Server (VCS) and Microsoft clusters.

Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for SFW installation and configuration information.

Supported software

Quick Recovery snapshots using Storage Foundation for Windows (SFW) require the following software:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW) with the FlashSnap option or Veritas Storage Foundation HA 5.1 Service Pack 2 for Windows (SFW HA) with the FlashSnap option.

Note: If using Quick Recovery in a cluster environment, Exchange must be installed on all cluster nodes.

- Microsoft Exchange servers and their operating systems:

Table 4-1 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2008 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2003 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2003 x64 Standard Edition or Enterprise Edition (SP2 required for all editions) ■ Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition (SP2 required for all editions)

Table 4-1 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2008 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

Storage requirements

The hardware for your SFW snapshot solution should include sufficient storage to be able to create dynamic volumes on separate disks for the following purposes:

- Databases
- Transaction logs
- Split-mirror snapshots of the database stores and transaction logs volumes

The system and boot volumes should reside on a separate disk (Harddisk0).

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the storage group, the snapshot set requires the same amount of space as the original volumes.

You can specify one or more snapshot sets for each volume with different disk locations for each. You can create more than one snapshot volume on the same disk as long as there is sufficient space available and as long as the snapshots are of different production volumes.

Disk groups must be of a Storage Foundation for Windows 4.0 or later version. It is necessary to upgrade any disk groups created using an earlier version of Volume Manager for Windows before implementing SFW snapshot solutions. Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a

Microsoft Disk Management Disk Group. For more information, see *Veritas Storage Foundation Administrator's Guide*

Methods of implementing Quick Recovery snapshots

Veritas Storage Foundation for Windows provides more than one method of implementing Quick Recovery snapshots.

[Table 4-2](#) summarizes the methods and when you would use them.

Table 4-2 Methods of implementing Quick Recovery snapshots

Task you want to accomplish	Method	For more information
<ul style="list-style-type: none"> ■ Set up and schedule multiple snapshot sets for multiple storage groups or ■ Update an existing snapshot set schedule 	<p>From the Solutions Configuration Center: Quick Recovery Configuration Wizard</p>	<p>“About the Quick Recovery Configuration Wizard” on page 41</p>
<ul style="list-style-type: none"> ■ Add a new snapshot set schedule for one storage group 	<p>From the VEA: VSS Exchange Snapshot Scheduler Wizard</p> <p>Alternatively, you can use the Quick Recovery Configuration Wizard.</p>	<p>“About the VSS Exchange Snapshot Scheduler Wizard” on page 42</p>
<ul style="list-style-type: none"> ■ Create a one-time snapshot as needed 	<p>From the VEA: VSS Exchange Snapshot Wizard and VSS Snapback Wizard</p> <p>Alternatively, you can use the Vxsnap utility from the CLI.</p>	<p>“About the VSS Snapshot and Snapback wizards and the vxsnap utility” on page 43</p>

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard provides the most complete method of configuring snapshot sets and is therefore recommended for initial configuration. The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirror for the snapshot set
- Creating the initial snapshot set by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy
- Periodically refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again

The wizard enables you to set up and schedule multiple snapshot sets for one or more storage groups of the selected Exchange instance. You can set up one or more schedules for each snapshot set.

You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. For example, you can schedule mirror preparation, the most time-consuming step, for a time of low activity.

The scheduling capability automates the process of refreshing the snapshot sets. At the scheduled times, the snapshot volumes are automatically reattached, resynchronized, and then split again. Once configured, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The wizard also enables you to save all configuration settings to a .tpl file (an XML file) that can be imported for use on other systems.

See [Chapter 6, “Implementing snapshot sets with the configuration wizard”](#) on page 69.

About the VSS Exchange Snapshot Scheduler Wizard

You can use the VSS Exchange Snapshot Scheduler Wizard to add a snapshot schedule for a selected storage group. This wizard uses the same scheduling process as the Quick Recovery wizard.

Unlike the Quick Recovery Configuration Wizard, the VSS Exchange Snapshot Scheduler Wizard does not prepare the snapshot mirror. To prepare the snapshot mirror you must use the VEA Prepare command (or the vxsnap command line utility) before running the VSS Exchange Snapshot Scheduler Wizard. Also, you can only schedule a snapshot for one storage group at a time.

You would typically use this wizard to add a schedule to the initial configuration that was created with the Quick Recovery Configuration Wizard. For example, you configured a daily snapshot for Quick Recovery use and now want to add a weekly snapshot on a different disk for backup use.

Note: Adding a snapshot schedule using the VSS Exchange Snapshot Scheduler does not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

See [Chapter 7, “Scheduling or creating an individual snapshot set”](#) on page 89.

About the VSS Snapshot and Snapback wizards and the vxsnap utility

If you need to create an individual one-time snapshot, you can do so either from the command line, using the vxsnap command line utility, or from the VEA console using the VSS Snapshot and Snapback wizards.

You must prepare the mirror for the snapshot manually, using the Prepare command. In addition, you cannot use the Snapshot wizard or command line to schedule periodic refreshing of the snapshot. Instead you manually reattach the mirror, allowing it to resynchronize, and then create a snapshot again from the resynchronized mirror.

Therefore, these methods are best suited for a one-time special-purpose snapshot. If you need to maintain a snapshot set for future use, you should instead set up a snapshot schedule using the VSS Exchange Snapshot Scheduler Wizard or the Quick Recovery Configuration Wizard.

See [Chapter 7, “Scheduling or creating an individual snapshot set”](#) on page 89. For Exchange 2007, if local continuous replication (LCR) or cluster continuous replication (CCR) is enabled, these SFW snapshot methods can also be used to create a snapshot of the passive copy of an Exchange 2007 storage group. The Quick Recovery Configuration Wizard does not support snapshots of the passive copy.

[Table 4-3](#) shows the corresponding vxsnap commands for the VSS wizards.

Table 4-3 Actions of VSS wizard and vxsnap command

Action	VSS wizard	vxsnap command
Creates a set consisting of snapshots of all the volumes in the specified Exchange storage group.	VSS Snapshot	<code>create</code>
Reattaches and resynchronizes a snapshot set to the production database volumes.	VSS Snapback	<code>reattach</code>

Note: The vxsnap commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode. For more information about the vxsnap command line utility, see [Chapter 10, “Vxsnap utility command line reference”](#) on page 169.

Planning your Quick Recovery solution

Consider the following questions as you plan your Quick Recovery solution:

- Should more than one snapshot set be created?
Symantec recommends maintaining a snapshot set that is refreshed after each full backup. A separate snapshot set that is refreshed after each incremental backup would reduce the time it takes to replay the logs in a point-of-failure recovery.
- How often and when will each snapshot set be refreshed?
- Is sufficient disk space available for multiple snapshot sets as well as the production volumes?
Each snapshot set contains split-mirror snapshots of all the volumes associated with a storage group and requires the same amount of space as the original volumes.
- Do you want to create a snapshot set schedule for use on multiple systems?
You can use the Quick Recovery Configuration Wizard to save a template to an XML file for use on other systems.

Backup types for snapshot sets

When creating a snapshot set, you can choose the backup type depending on how you plan to use the snapshot set:

- Copy backup is generally used for a snapshot set to be maintained on-host and used for quick recovery of a missing or corrupt Exchange storage group or database. This mode creates a snapshot or copy of the transaction log, but does not truncate it. You can choose whether or not to run the Eseutil consistency check.
- Full backup is generally used for a snapshot set to be used for snapshot-assisted backup. Full backup runs the Eseutil consistency check and truncates the transaction logs.

About logs

Microsoft Exchange employs “write-ahead” logging to improve performance. As transactions occur, they are appended to a transaction log. Transactions are committed to the database when resources permit. A checkpoint file (EOn.CHK) tracks the location of the last transaction in the log files that has been successfully committed to the database. Transaction logs can be used to roll forward a database to achieve a point-of-failure recovery.

Note: Do not enable the circular logging option. If circular logging is enabled, you will not be able to 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.

Recommendations and best practices

You should review the recommendations and best practices for the following:

- [“Recommendations for Exchange storage configuration”](#) on page 46
- [“Recommendations for maintenance and backups”](#) on page 46
- [“VCS cluster considerations”](#) on page 47
- [“Microsoft cluster considerations”](#) on page 48
- [“VVR considerations”](#) on page 50
- [“Exchange 2007 passive copy snapshot considerations”](#)

Recommendations for Exchange storage configuration

To use the Quick Recovery snapshot functionality with Exchange storage groups, you must place the storage groups on Veritas Storage Foundation for Windows (SFW) dynamic volumes. The following recommendations enable you to take advantage of SFW storage configuration functionality as you manage your Exchange storage:

- Database stores and transaction logs for each storage group must be stored on disks contained within a single dynamic disk group.
- Each database should be in a separate volume, but the volumes may share the same dynamic disks.
- Mailbox stores and public stores must be stored on separate volumes in order to be able to recover each independently.
- Database stores and transaction logs must be in separate volumes in order to perform a roll-forward recovery to the point of failure.
- Database stores and transaction logs should be on separate disks so that disk failure does not affect both the database stores and transaction logs.
- Transaction logs should always be configured in a redundant layout. The preferred software layout is RAID 0+1 (mirrored striped) volumes as this provides better read and write performance than RAID 1 (mirrored) alone. The transaction log will generate the most I/O and thus should use the highest performance disks available.
- The preferred layout for the database stores is hardware RAID 5, software RAID 1 (mirrored with logging enabled) or software RAID 0+1 (mirrored striped).

Note: FlashSnap is not supported for software RAID 5 volumes.

- By default, a First Storage Group is created in the install location, which is the boot drive. Snapshots cannot be created on the boot drive. In order to use VSS snapshots on the storage group, you must move the storage group components from the boot drive to an SFW dynamic volume on another drive.
See [“Moving the First Storage Group \(Exchange 2003\)”](#) on page 62.
See [“Moving the First Storage Group \(Exchange 2007\)”](#) on page 67.

Recommendations for maintenance and backups

- Locate the snapshot volumes for each storage group on separate disks from snapshots of other storage groups. This is recommended so that the process

of creating the snapshot of one storage group doesn't interfere with any operations on another storage group.

- Each snapshot set is associated with a metadata XML file that is needed for Quick Recovery restoration. Create a unique name for the metadata XML file of each snapshot set you create and maintain. The snapshot XML files should be stored separately from the volumes that are included in snapshots. Snapshot schedules have additional XML files associated with them. By default, the snapshot XML files are created in the following paths.

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application
Data\VERITAS\VSSXML\application name
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\VSSXML\application name
```

- For Quick Recovery solutions, Symantec recommends that you create or refresh a snapshot set immediately after a Full Backup just after the database has been checked for corruption and the transaction logs have been truncated. Thus, you are assured an image of a clean database. Additionally, you may wish to create another snapshot set after an Incremental Backup. Create this snapshot set on a separate set of disks rather than refreshing the snapshot set taken after the Full Backup. This practice ensures you are not overwriting a snapshot set of a clean database with an image of a potentially corrupted database.

VCS cluster considerations

In a VCS cluster environment, observe the following precautions:

- The XML metadata file for each snapshot set along with files that store snapshot schedule information are created in a folder on the local drive by default. In a cluster environment, store these files on shared storage so that the files are available from all nodes in the cluster. The snapshot XML files should be stored separately from the volumes that are included in snapshots.
 - If you use the Quick Recovery Configuration Wizard to create the snapshot set, you can use the wizard to specify the file path to the appropriate volume.
 - If you use a VSS wizard to create the snapshot set, you can store the XML files in a location of your choice using one of the following methods:
 - Edit the directory path in the Directory field in the VSS wizard.
 - Use a text editor to create a text file named **redirect.txt**. This text file should contain a single text line specifying the full path to

the location of the metadata file, for example, `G:\BackupSets`. Save the `redirect.txt` file in the default VSS XML file directory `C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML` on each node of the cluster.

- When using `vxsnap` utility commands that require the filename attribute, specify the full path to the location of the XML metadata file.
- If you plan to use the Quick Recovery or VSS Snapshot Scheduler wizard to specify scripts (commands) to be run before or after a snapshot, store the scripts on shared storage so that they are available to all nodes.
- If you set up a snapshot schedule with the Quick Recovery wizard and later add a node to the cluster, you can run the wizard again to synchronize schedules on the existing nodes with the new node.
- If you set up a snapshot schedule with the VSS Snapshot Scheduler wizard, before adding or removing a node from a VCS cluster setup, delete the schedules and then recreate the schedules on the required node.
- Exchange must be installed on all cluster nodes.

Microsoft cluster considerations

In a Microsoft clustering environment (MSCS or failover clustering), observe the following precautions:

- Store XML files required for snapshot operations on shared storage. The XML metadata file for each snapshot set along with files that store snapshot schedule information are created in a folder on the local drive by default. In a cluster environment, store these files on shared storage so that the files are available from all nodes in the cluster. The snapshot XML files should be stored separately from the volumes that are included in snapshots.
- If you use the Quick Recovery Configuration Wizard to create the snapshot set, you can use the wizard to specify the file path to the appropriate volume.
- If you use a VSS wizard to create the snapshot set, you can store the XML files in a location of your choice using one of the following methods:
 - Edit the directory path in the Directory field in the VSS wizard.
 - Use a text editor to create a text file named "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, `G:\BackupSets`. Save the **redirect.txt** file in the default VSS XML file directory

C:\Program Files\Veritas\Veritas Volume Manager
 5.1\VSSXML on each node of the cluster.

- When using the vxsnap utility, specify the full path to the location of the XML metadata file in commands that require the filename attribute.
- If you plan to use the Quick Recovery or VSS Snapshot Scheduler wizard to specify scripts to be run before or after a snapshot, store the scripts on shared storage so that they are available to all nodes.
- If you set up a snapshot schedule with the Quick Recovery wizard and later add a node to the cluster, you can run the wizard again to synchronize schedules on the existing nodes with the new node.
- Exchange must be installed on all cluster nodes.
- If multiple Exchange virtual servers are online on the same server, ensure that they do not contain storage groups with the same name. Otherwise, snapshot operations on the components with the same name may fail. Rename any storage groups that have the same name, as follows:
 - For Exchange 2003, in Exchange System Manager, right-click the storage group that you want to rename and click Rename.
 - For Exchange 2007, in the Exchange Management Console, right-click the storage group that you want to rename and click Properties. In General Properties, change the name in the editable box and click Apply.
- Before performing either a point-in-time recovery or a roll-forward recovery to the point of failure of an Exchange storage group, use the Microsoft clustering software console to offline the following resources:
 - Exchange HTTP Virtual Server Instance 100
 - Exchange IMAP4
 - Exchange Message Transfer Agent Instance (Exchange 2003 only)
 - Exchange POP3
 - Exchange Routing Service Instance
 - Exchange SMTP Virtual Service Instance

The Exchange Information Store, Exchange SA, and VMDg resources in the Exchange resource group must remain online as these resources control the Exchange writer and enable the VSS snapshot and restore processes.
- Pause all passive nodes in the cluster before beginning a roll-forward recovery to the point of failure. This is necessary to prevent inadvertent failover during the recovery process. A failover can occur when the

replaying of a large number of logs prevents Exchange from responding to the MSCS IsAlive call.

- For Exchange 2007: Before performing a restore operation on a passive copy snapshot, manually dismount the databases and set for overwrite by restore.

VVR considerations

In a VVR environment, observe the following precautions:

- Store the XML metadata file and other snapshot related files on a volume that is included in the replicated data set so that the metadata file is available at the secondary site. Additionally, if VCS is used, store the metadata file on a volume in the cluster disk group associated with the Exchange storage group so that the metadata file is available from all nodes in the cluster.

See “[VCS cluster considerations](#)” on page 47.

- During a point-in-time recovery, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

See “[Post-recovery steps](#)” on page 123 in [Chapter 8, “Recovering Exchange mailbox storage group or databases”](#).

Exchange 2007 passive copy snapshot considerations

Exchange 2007 supports Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR). When either of these are enabled, the Microsoft Exchange Writer Replica instance of the VSS Writer is enabled.

There are limitations on SFW support for snapshots of the passive (replicated) copy of the storage group. The Quick Recovery Wizard and the VSS Scheduler Wizard do not support snapshots for the passive copy. Therefore, you cannot schedule snapshots of the passive copy.

However, SFW does support creating snapshots of the LCR or CCR passive storage group with the VSS Exchange Snapshot Wizard or with the vxsnap CLI command. Creating a snapshot of a passive storage group requires using the Microsoft Exchange Writer Replica instance to prepare the snapshot mirror and to create the snapshot. If replication is enabled:

- The Microsoft Exchange Writer Replica instance is displayed in the tree view of the VEA under the Applications node. Right-clicking the Exchange

Replica node displays a context menu that shows VSS Snapshot, VSS Snapback, and VSS Refresh selections.

- Using the `vxsnap` CLI command, you can specify Microsoft Exchange Writer Replica (instead of Microsoft Exchange Writer) for the `writer` option in the required commands.

Using the passive copy for snapshots is described further in [Chapter 7, “Scheduling or creating an individual snapshot set”](#).

Although you cannot restore a snapshot to the passive copy, you can use the passive copy snapshot to restore a storage group or database after first manually performing a "Restore-StorageGroupCopy" on the storage group. However, note the following considerations:

- The LCR copy is automatically disabled when running the `Restore-StorageGroupCopy` cmdlet.
- If a snapshot schedule exists for the active storage group, running the `Restore-StorageGroupCopy` cmdlet makes the schedule invalid, and you must delete the invalid schedule before running the restore operation. The cmdlet changes which volumes are the active copy. The previous passive volumes are now active, and vice versa. The scheduled snapshots point the active writer instance, Microsoft Exchange Writer, toward what are now passive volumes, so the snapshot operations fail. The schedule must be deleted so that the active writer can be pointed toward the new active volumes in order to do the restore operation.

See [“Recovery using an Exchange 2007 passive copy snapshot”](#) on page 113 in [Chapter 8, “Recovering Exchange mailbox storage group or databases”](#).

Configuring Exchange for Quick Recovery snapshots

This chapter includes the following topics:

- [Tasks for configuring Exchange storage](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Creating dynamic disk groups](#)
- [Creating dynamic volumes](#)
- [Pointing the databases and log paths to the SFW volumes \(Exchange 2003\)](#)
- [Moving the First Storage Group \(Exchange 2003\)](#)
- [Pointing the databases and log paths to the SFW volumes \(Exchange 2007\)](#)
- [Moving the First Storage Group \(Exchange 2007\)](#)

Tasks for configuring Exchange storage

[Table 5-1](#) outlines the high-level objectives and the tasks to complete each objective for configuring Exchange storage for Quick Recovery snapshots.

Table 5-1 Tasks for configuring Exchange storage for Quick Recovery snapshots

Objective	Tasks
“Reviewing the prerequisites” on page 54	<ul style="list-style-type: none"> Verifying hardware and software prerequisites
“Reviewing the configuration” on page 55	<ul style="list-style-type: none"> Reviewing the sample configuration
“Creating dynamic disk groups” on page 56	<ul style="list-style-type: none"> Creating dynamic disk group for each Exchange storage group
“Creating dynamic volumes” on page 57	<ul style="list-style-type: none"> Creating volumes for the database stores and transaction logs in each storage group
“Pointing the databases and log paths to the SFW volumes (Exchange 2003)” on page 60	<ul style="list-style-type: none"> (Exchange 2003) Using Exchange to set the database and log paths for new or existing storage groups and databases to point to the SFW volumes
“Moving the First Storage Group (Exchange 2003)” on page 62	<ul style="list-style-type: none"> (Exchange 2003) Moving the First Storage Group from its default location on the boot volume
“Pointing the databases and log paths to the SFW volumes (Exchange 2007)” on page 65	<ul style="list-style-type: none"> (Exchange 2007) Using Exchange to set the database and log paths for new or existing storage groups and databases to point to the SFW volumes
“Moving the First Storage Group (Exchange 2007)” on page 67	<ul style="list-style-type: none"> (Exchange 2007) Moving the First Storage Group from its default location on the boot volume

Reviewing the prerequisites

This solution assumes that the required SFW software is already installed and configured.

Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation and configuration information.

The solution also assumes that the required Exchange software is installed.

You should also familiarize yourself with the system requirements and recommendations and best practices for configuring Exchange storage for snapshot solutions.

See [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).

Reviewing the configuration

[Table 5-3](#) shows the object names that are used to illustrate the tasks you perform when configuring Exchange storage with Veritas Storage Foundation for Windows:

Table 5-3 Object names for Exchange

Name	Drive Letter	Object
Exch1		Exchange Server host name
SG2		Exchange Storage Group (Component Name)
Exch1-SG2		Dynamic disk group associated with the SG2 storage group
Harddisk1, Harddisk2, Harddisk3		Disks included in Exch1-SG2
SG2-tlogs	I:	Transaction logs volume
SG2-DB1	G:	Mailbox store database volume
SG2-DB2	H:	Mailbox store database volume
SG2-tlogssnap		Snapshot volume of SG2-tlogs
SG2-DB1snap		Snapshot volume of SG2-DB1
SG2-DB2snap		Snapshot volume of SG2-DB2
Image1.xml		Name for the metadata file that contains information about the snapshot set

Creating dynamic disk groups

Create one dynamic disk group for each Exchange storage group.

If your Exchange production server is in a clustered environment, choose the cluster disk group option.

See *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange* for further configuration information on the SFW HA clustered environment.

Note: Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before implementing SFW snapshot solutions. Quick Recovery snapshots are supported only on volumes belonging to an SFW dynamic disk group. They are not supported on volumes belonging to a Microsoft Disk Management Disk Group. For more information, see *Veritas Storage Foundation Administrator's Guide*.

To create a dynamic disk group from the VEA console

- 1 Click **Start > Programs > Symantec > Veritas Storage Foundation>Veritas Enterprise Administrator** and if prompted to select a profile, select a profile (or Default).
- 2 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
For the local system you can specify **localhost**.
- 3 If prompted to do so, specify the user name, password, and domain for the system.
- 4 In the tree expand the system name and expand the storage agent.
- 5 Right-click **Disk Groups**, and click **New Dynamic Disk Group**.
- 6 On the **Welcome** screen of the New Dynamic Disk Group Wizard, click **Next**.
- 7 Enter a name for the disk group (for example, Exch1-SG2).
- 8 For an off-host or cluster environment, choose from the following:
 - For a cluster environment, check the **Create cluster group** check box.
 - For an off-host environment, check the **Add private group protection** check box.
- 9 Select the appropriate disks in the Available disks list and click the **Add** button to move them to the Selected disks list.
- 10 Click **Next**.

- 11 Click **Next** to upgrade the selected disks.
- 12 Click **Finish** to create the new disk group.

To create a dynamic disk group from the command line

- Type the following command:

For an on-host environment

```
vxdg -gExch1-SG2 init Harddisk1 Harddisk2 Harddisk3
```

For an off-host environment

```
vxdg -gExch1-SG2 -R init Harddisk1 Harddisk2  
Harddisk3
```

For a cluster environment

```
vxdg -gExch1-SG2 -s init Harddisk1 Harddisk2  
Harddisk3
```

where Exch1-SG2 is the name of the dynamic disk group you want to create and Harddisk1, Harddisk2, and Harddisk3 are the disks included in the dynamic disk group.

For the complete syntax of the `vxdg init` command, see the *Veritas Storage Foundation Administrator's Guide*.

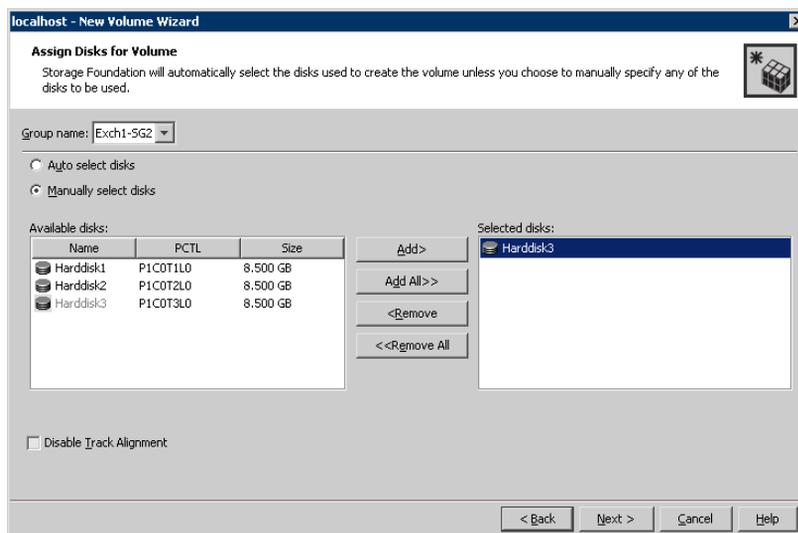
Creating dynamic volumes

For each storage group, create a volume for the log files and an additional volume for each database.

To create a dynamic volume from the VEA console

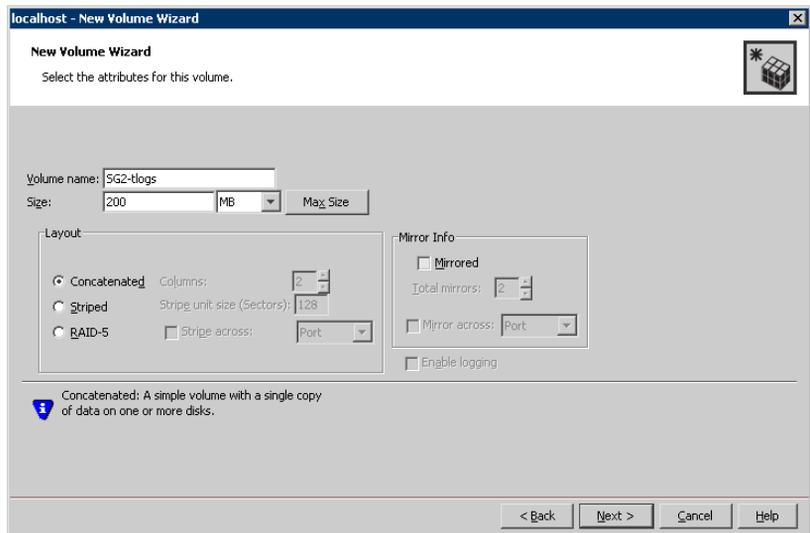
- 1 Start the VEA and connect to the appropriate host.
- 2 In the tree, expand the system name, expand the storage agent, and then expand **Disk Groups**.
- 3 Right-click on the disk group in which to create the volumes (for example, Exch1-SG2), and click **New Volume**.
- 4 In the Welcome panel of the New Volume Wizard, click **Next**.

5 Select the disks for the volume:



- Confirm that the **Group name** is the correct disk group (Exch1-SG2). If necessary, select the correct disk group name from the drop-down menu.
- Specify automatic or manual disk selection. Symantec recommends using the **Manually select disks** option.
- Select the appropriate disks in the Available disks list, and click the **Add** button to move them to the Selected disks list.
- You may also check Disable Track Alignment to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

6 Specify the parameters of the volume:



- Enter the volume name (for example, SG2-tlogs).
 - Enter the size.
 - Select the layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 7** Assign a drive letter to the volume (for example, I: for the SG2-tlogs volume) and click **Next**.
- 8** Create an NTFS file system:
- Accept the default **Format this volume**.
 - Click **NTFS**.
 - Select an allocation size.
 - Accept the default file system label, which is the same as the volume name you entered previously or enter a file system label.
 - If desired, select **Perform a quick format**.
 - Click **Next**.
- 9** Review the volume specifications, then click **Finish** to create the new volume.
- 10** Repeat the previous steps as necessary to create volumes for the databases in the storage group. For example, create volumes SG2-DB1 and SG2-DB2.

To create a volume from the command line

- 1 Type the CLI command, as in the following example:

```
vxassist [-b] -gExch1-SG2 make SG2-tlogs 200 DriveLetter=I
```

This example command will create a 200 MB volume named SG2-tlogs on drive letter I: in the Exch1-SG2 storage group.

Note: This command does not format the volume (a file system is not created on the volume). You must use the operating system format command to format the volume.

- 2 Modify and repeat this command as necessary to create volumes for the databases in the storage group.

For the complete syntax of the `vxassist make` command, see *Veritas Storage Foundation Administrator's Guide*.

Pointing the databases and log paths to the SFW volumes (Exchange 2003)

After you create volumes in Veritas Storage Foundation for Windows (SFW), you set the Exchange 2003 log file paths, system file paths, and database paths to point to the SFW volumes.

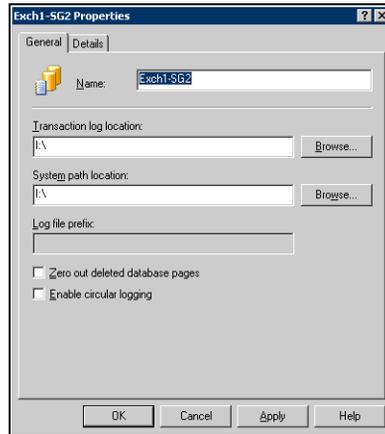
For Exchange 2003, set the path for the transaction log location and system path location fields to point to the log volume. Set the path for the Exchange database and Exchange streaming database files to point to the appropriate database volumes. You use the Exchange System Manager to specify the locations.

Optional steps for creating a new storage group and mailbox stores are included in the procedure.

To point the Exchange 2003 databases and log paths to the SFW volumes

- 1 Click **Start > All Programs > Microsoft Exchange > System Manager** to open the Exchange 2003 System Manager.
- 2 In the appropriate Administrative Group, expand **Servers** and expand the appropriate Exchange server.
- 3 Choose one of the following:
 - Right-click an existing storage group, for example SG2, and click **Properties**.
 - To create a new storage group, right-click the appropriate Exchange server, click **New** and click **Storage Group**.

4 Complete the following on the Properties sheet and click **OK**:



Name	For a new storage group, enter the name storage group name (for example, SG2).
Transaction log location	Click the Browse button and set the transaction log path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \.
System path location	Click the Browse button and set the path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \.
	The paths for the Transaction log location and System path location should be the same.
Enable circular logging	Make sure that the Enable circular logging check box is not checked.

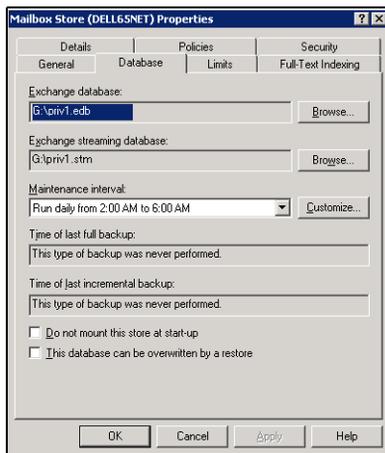
5 Right-click on the storage group and choose one of the following:

- For an existing storage group, right-click on a mailbox store and click **Properties**.
- For a new storage group, click **New** and click **Mailbox Store**.

6 Choose one of the following:

- For an existing storage group, proceed to [step 7](#).
- For a new storage group, in the **General** tab of the Properties sheet enter the name of the new mailbox store (for example, SG2-DB1).

- 7 Click the **Database** tab, set the paths for the .edb and .stm files for the database as follows, and click **OK**:



Exchange database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing EDB files (for example, G: \).

Exchange streaming database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing STM files (for example, G: \).

- 8 Click **Yes** to mount the store.
- 9 Repeat [step 5](#) through [step 8](#) to create or set the paths for other mailbox stores. For example, create another mailbox store mounted on the H: drive, SG2-DB2.

Moving the First Storage Group (Exchange 2003)

When Exchange Server is installed, the First Storage Group is created in the installation location, which by default is the boot drive. The boot drive cannot be snapshotted. In order to implement a snapshot solution you must move the components (mailbox store and public folder store) of the First Storage Group to new volumes not on the boot drive.

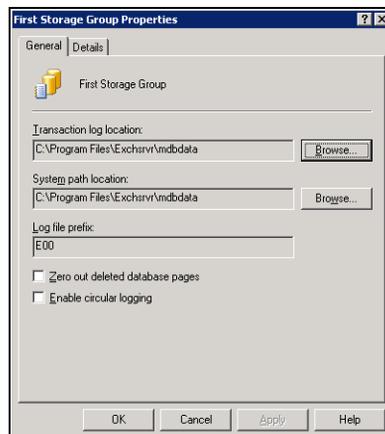
This involves the following procedures:

- Create three new volumes, one each for the mailbox store, public folder stores, and transaction log
- Move the transaction log to the new volume
- Move the mailbox and public stores to the new volumes

See the Microsoft Exchange 2003 documentation for additional information on creating new storage groups and databases.

To move the First Storage Group transaction log to a new volume

- 1 Click **Start > All Programs > Microsoft Exchange > System Manager** to open the System Manager.
- 2 In the appropriate Administrative Group, expand **Servers** and expand the appropriate Exchange server.
- 3 Right-click **First Storage Group** and click **Properties**.
- 4 On the Properties sheet, click the **General** tab, and point to the new volume for the transaction log as follows, and click **OK**:



Transaction log location Click the **Browse** button and set the transaction log path to the drive letter or mount point of the volume created for the transaction log.

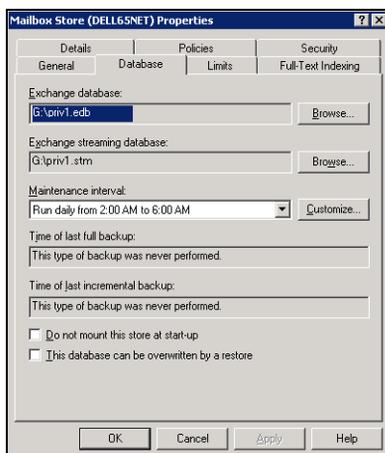
System path location Click the **Browse** button and set the path to the drive letter or mount point of the volume created for the transaction log of the storage group. The paths for the Transaction log location and System path location should be the same.

Enable circular logging

Make sure that the Enable circular logging check box is not checked.

To move the First Storage Group mailbox and public stores to new volumes

- 1 In the System Manager (**Start > All Programs > Microsoft Exchange > System Manager**), under the appropriate Exchange server and First Storage Group, right-click on the **Mailbox Store** and select **Properties**.
- 2 On the Properties sheet, click the **Database** tab and set the paths for the EDB and STM files for the database:



Exchange database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing EDB files (for example, G: \).

Exchange streaming database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing STM files (for example, G: \).

- 3 Click **Yes** to mount the store.
- 4 Under the First Storage Group, right-click on **Public Store** and select **Properties**.
- 5 Repeat [step 2](#) and [step 3](#) to move the public folder store to the new volume created for it.

Pointing the databases and log paths to the SFW volumes (Exchange 2007)

After you create the volumes in SFW, you must set the Exchange 2007 log file paths, system file paths, and database paths to point to the SFW volumes.

For each Exchange 2007 storage group, set the log files path and system files path to point to the log volume. For each database, set the database path to point to the appropriate database volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new storage groups or databases or change the locations for existing storage groups or databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

The following procedures can be used to change paths for existing storage groups and databases.

See the Microsoft Exchange 2007 documentation for additional information on creating new storage groups and databases.

To use the Exchange Management Console to change log file and system file locations for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the storage group for which you want to change the log file location.
- 4 In the work pane, right-click the storage group for which you want to change the log file location and click **Move Storage Group Path**. The Move Storage Group Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the log files, and create a folder for the log files. Repeat for the system files. The volume should be the same for both.
- 6 Click **Move**. A warning appears that all databases in the storage group must be temporarily dismounted, which will make them inaccessible to any user. To continue, click **Yes**.

- 7 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Storage Group Path wizard.

To use the Exchange Management Console to change the database file location for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the database for which you want to change the database file location.
- 4 In the work pane, right-click the desired database and click **Move Database Path**. The Move Database Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the database file, and create a folder for it as needed.
- 6 Click **Move**. A warning appears indicating that the database must be temporarily dismounted, which will make it inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the database files were moved successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Moving the First Storage Group (Exchange 2007)

When Exchange Server is installed and the First Storage Group is created, the First Storage Group is mapped to the boot drive by default. The boot drive cannot be snapshotted. In order to implement a snapshot solution you must move the components of the First Storage Group to new volumes not on the boot drive. This involves the following procedures:

- Create two new volumes, one for the mailbox store and the other for the transaction logs and system files.
See “[Creating dynamic volumes](#)” on page 57 for instructions.
- Move the log, system, and database files to the appropriate new volume
See “[Pointing the databases and log paths to the SFW volumes \(Exchange 2007\)](#)” on page 65.

Implementing snapshot sets with the configuration wizard

This chapter covers the following topics:

- [About the Quick Recovery Configuration Wizard](#)
- [Tasks for implementing snapshot sets with the configuration wizard](#)
- [Reviewing the prerequisites](#)
- [Scheduling and creating snapshot sets](#)
- [Viewing the status of scheduled snapshots](#)
- [Troubleshooting scheduled snapshots](#)
- [Deleting or modifying schedules](#)
- [Synchronizing schedules after adding a cluster node](#)

About the Quick Recovery Configuration Wizard

The Quick Recovery Configuration Wizard enables you to schedule all phases of snapshot creation and maintenance:

- Preparing the mirrors
- Creating the initial snapshot sets
- Periodically refreshing the snapshot sets

The wizard enables you to set up and schedule multiple snapshot sets for one or more Exchange components. For Microsoft Exchange 2003 and 2007, Exchange

components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

You can set up different schedules for each snapshot set. You can schedule when to prepare the snapshot mirrors, when to create the initial snapshot sets, and when to refresh the snapshot sets, enabling you to establish a schedule that best suits your site. The scheduling capability automates the process of refreshing the snapshot sets. At the scheduled time, the snapshot volumes are automatically reattached, resynchronized, and then split again.

The snapshot creation process integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes. The snapshot is done while the databases are online and without disrupting the email flow.

Once configured and applied, snapshot schedules are maintained by the Veritas Scheduler Service, which runs in the background.

The wizard includes the following settings:

- Which Exchange component to snapshot.
- Number of snapshot sets for each Exchange component
- Volume layout and disk location for each snapshot set
- When to prepare the selected volumes for the snapshots
- When and how often snapshots occur
You can schedule snapshot run days using days of the week, days of the month, and include or exclude dates. You can specify that a schedule recurs daily or uses some other recurrence interval. You can also schedule a daily time window in which the snapshots occur, and the intervals for refreshing snapshot sets within the time window.
- Backup type (Full, Copy, or Copy plus Eseutil)

Optionally, you can also specify scripts to run before and after snapshots.

After you configure the settings, you can do one or both of the following:

- Save the settings in a template file for reuse on other systems. Even if you do not save the settings to a template file, they are still stored for the Exchange component to which they apply. You can access them later by selecting the same instance and Exchange component in the wizard.
- Implement the settings immediately.

About snapshot sets

Because a snapshot set contains a split-mirror snapshot copy of each of the volumes in the Exchange component, the snapshot set requires the same amount of space as the original volumes. For Microsoft Exchange 2003 and

2007, Exchange components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

Each snapshot set can be created as a different backup type: full copy including Eseutil consistency check and transaction log truncation, copy with Eseutil, or copy without Eseutil or transaction log truncation.

Snapshot set creation has the following stages, all of which you can schedule using the Quick Recovery Configuration Wizard:

- Preparing the mirror for the snapshot set. This stage takes a while and should be scheduled for a time of low activity. The remaining stages take little time.
- Creating the initial snapshot by splitting the mirror so that it is no longer synchronized with the original volume and becomes a point-in-time copy.
- Refreshing (resynchronizing) the split-mirror snapshot with the original volume, and then splitting the mirror again. This stage occurs for each snapshot set at the times or frequency that you specify. For example, you can schedule one snapshot to occur every 30 minutes starting at 9 A.M. and ending at 6 P.M.

About snapshot templates

A snapshot template contains all the settings that you implemented with the Quick Recovery Configuration Wizard for the selected Exchange component. For Microsoft Exchange 2003 and 2007, Exchange components are defined as storage groups. For Microsoft Exchange 2010, Exchange components are defined as databases.

After configuring settings with the wizard, you can save the template settings to a template (.tpl) file (an XML file) for use on another Exchange component. This enables you to re-use the same schedule or use it with minor modifications. If you apply the settings without saving, they are stored in the registry so that you can view or modify them in the wizard, but they are not available to import for another instance or Exchange component.

Before you can use the template settings on another system, you must first copy the template (.tpl) file to that system. You can then import the file while using the Quick Recovery Configuration Wizard.

If settings already exist for an Exchange component that you selected in the wizard, you cannot import a template to overwrite those settings. You can only modify the settings in the wizard. However, you can import a template for another Exchange component for which settings do not exist.

Templates store the following information:

- The selected Exchange component

- The number of snapshot sets for each Exchange component
- When to prepare the snapshot mirrors
- The snapshot schedule and rules related to the schedule, including backup type and the names of optional scripts
- The current date and time when the template is created
- The application name

Templates and multiple components

When you apply a template that has multiple components, the wizard first attempts to match the names of components (storage group for Microsoft Exchange 2003 and 2007, or database for Microsoft Exchange 2010) in the template to the Exchange component you selected in the wizard.

If it matches a name, it applies the information stored under that name in the template to that Exchange component. For example, if you select Exch_SG1 for a Microsoft Exchange 2003 or 2007 component or Exch1-DB1 for a Microsoft Exchange 2010 component in the wizard and the template has settings for these Exchange components, it applies those settings.

An Exchange component selected in the wizard may not match the names of any Exchange component in the template. In that case, the wizard applies the information for the first unapplied component in the template to the first selected Exchange component that does not match any name in the template. It continues in that sequence.

If you selected more Exchange components in the wizard than the number of components in the template, the wizard prompts you to fill in any required information for the remaining Exchange components.

Templates and schedule start dates

Templates contain settings for one or more schedule start dates, including a mirror preparation date, a "Start on" date, and a "Schedule takes effect on" date.

If you import a template after a start date has elapsed, the wizard tries to maintain the same delta between the template import date and a start date as the delta between the template creation date and a start date. It shifts the dates forward to maintain this delta.

Therefore, after importing the template, use the wizard to review the settings and adjust dates as necessary to meet your needs.

If you import a template that uses the current date as the "takes effect on" date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, a schedule includes two daily snapshots, one in the morning and one at night, and you import the

schedule in the afternoon. In this case neither of the snapshots will occur on the current date. Instead the effective date is shifted ahead one day.

Tasks for implementing snapshot sets with the configuration wizard

[Table 6-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 6-1 Tasks for implementing snapshot sets with the configuration wizard

Objective	Tasks
“Reviewing the prerequisites” on page 73	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Scheduling and creating snapshot sets” on page 74	<ul style="list-style-type: none"> ■ Running the Quick Recovery Configuration Wizard to configure and schedule snapshot sets ■ Optionally, saving the configuration and schedules to a template file ■ Applying the configuration
“Viewing the status of scheduled snapshots” on page 83	<ul style="list-style-type: none"> ■ Viewing the status of scheduled snapshots to determine whether they were successful
“Deleting or modifying schedules” on page 86	<ul style="list-style-type: none"> ■ Deleting schedules that are no longer needed or modifying schedules.
“Synchronizing schedules after adding a cluster node” on page 87	<ul style="list-style-type: none"> ■ In a clustered environment, synchronizing schedules after adding a cluster node

Reviewing the prerequisites

Before running the Quick Recovery Configuration Wizard, you should do the following:

- Ensure that your system hardware and software meets the requirements. See [“System requirements”](#) in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).

- Review the recommendations and best practices.
See [“Recommendations and best practices”](#) on page 45 in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).
- Set up your Exchange component (storage group or database depending upon the version of Microsoft Exchange you are using) for use with Storage Foundation for Windows (SFW).
See [Chapter 5, “Configuring Exchange for Quick Recovery snapshots”](#) on page 53.
- For Microsoft Exchange 2003 or 2007, ensure that the mailbox stores for which you plan to schedule snapshot operations are mounted. For Microsoft Exchange 2010, ensure that the mailbox databases for which you plan to schedule snapshot operations are mounted.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- Ensure that you are logged in as a domain administrator or as a member of the Domain Admins group.
- If a firewall exists between the wizard and any systems it needs access to, set the firewall to allow both ingoing and outgoing TCP requests on port 7419.
- In a clustered server environment, ensure that the Veritas Scheduler Service is configured and running with domain administrator privileges on all nodes in the cluster.
- To use a previously created template, copy it to the server that you are configuring.
- Ensure that the Microsoft Software Shadow Copy Provider service is running. For a clustered server environment, ensure that the service is running on all systems in the cluster.

Scheduling and creating snapshot sets

You schedule and create snapshot sets using the Quick Recovery Configuration Wizard. You can also use the wizard to modify or delete existing schedules or import a schedule template.

See [“About the Quick Recovery Configuration Wizard”](#) on page 69.

You should ensure that you meet the prerequisites before you begin using the wizard.

See [“Reviewing the prerequisites”](#) on page 73.

In addition, if you plan to reuse settings for the Exchange component (storage group or database depending upon the version of Microsoft Exchange you are using) on other instances, you should be familiar with the information about snapshot templates.

See “[About snapshot templates](#)” on page 71.

Note: For Exchange 2007, if using local continuous replication (LCR), you can schedule snapshots only for the active copy, not the passive copy.

Note: For Exchange 2010, if using replication in a DAG, you can schedule snapshots only for the active copy, not the passive copy.

To schedule and create snapshot sets

- 1 Start the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**), expand the application on the Solutions tab, expand **Quick Recovery Configuration**, and start the Quick Recovery Configuration Wizard.
- 2 In the Welcome panel, review the information to ensure you meet the requirements and click **Next**.
- 3 In the System Selection panel, specify the fully qualified host name or IP address of the system that is running the application (or specify localhost for the local system), and click **Next**.
 For failover of an application virtual server in a clustered server (high availability) environment, Symantec recommends that you specify the virtual server name or virtual IP address.
 Ensure that the disk groups for the application are imported to the selected system.
- 4 In the Instance Selection panel, specify the following options and click **Next**:

Set up Quick Recovery for	For Exchange 2003 or 2007, select the Exchange virtual server name. For Exchange 2010, select the system name of the Exchange mailbox server. The selection details are displayed in the panel.
---------------------------	--

Select objects or view details Depending upon your version of Microsoft Exchange, the Components list shows either storage groups or databases. Select the component for which you want to configure the snapshot sets.

For Exchange 2010, you can select either VCS clustered databases or standalone, but not both together. If you are selecting clustered databases, they must be configured to fail over to the same list of systems.

The wizard checks for existing schedules for the selected component. The wizard cleans up any obsolete schedule files and entries. In a clustered server environment, the wizard will synchronize schedules between cluster nodes to ensure that all existing schedules are available on the selected node.

- 5 The wizard validates that the volumes containing the schedule configuration files are mounted. If a volume is not mounted, the Mount Details panel displays the information about the missing drive. Mount the missing drive and click **Refresh**, then click **Next**. Otherwise, the schedule is deleted.
- 6 For existing schedules, the Synchronizing Schedules panel displays the status of schedule synchronization and cleanup. If schedule synchronization fails in a clustered environment, restart any cluster nodes that are down and then restart the wizard. Otherwise, click **Next**.
- 7 In the Template Selection panel, select from the following options and click **Next**:

Create or modify a template Create or modify the Quick Recovery settings for the selected components.

If at least one of the selected components has existing settings, this option is the only choice. In this case, the existing settings are displayed for you to view or modify as you continue through the wizard panels.

If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.

Import a template Import Quick Recovery settings from a template (.tpl) file created by the Quick Recovery wizard earlier. Browse to the file location and select the file. The file must be located on the system for which Quick Recovery is being configured. The file also must be of the same application type as the application you are configuring.

This option is available only if no settings exist for a selected component.

- 8 In the Number of Snapshot Sets panel, select how many sets of snapshots to create for each database. Remember that each snapshot set uses the same amount of space as the original volumes.
- 9 In the Snapshot Volume Assignment panel, expand the snapshot set and make the following selections:

Snapshot set In the box under the Snapshot Set column heading, optionally edit the snapshot set name. If you edit the name, ensure that the name is unique among all snapshot sets for all databases.

If you are modifying an existing schedule, the snapshot set name cannot be changed.

XML Metadata File Name Specify a name for the XML file that will be associated with the snapshot set. This file is used for recovery, so you may want to assign a name to easily identify it for that purpose. Ensure that the XML file for each snapshot set has a unique name.

If you are modifying an existing schedule, the XML file name cannot be changed.

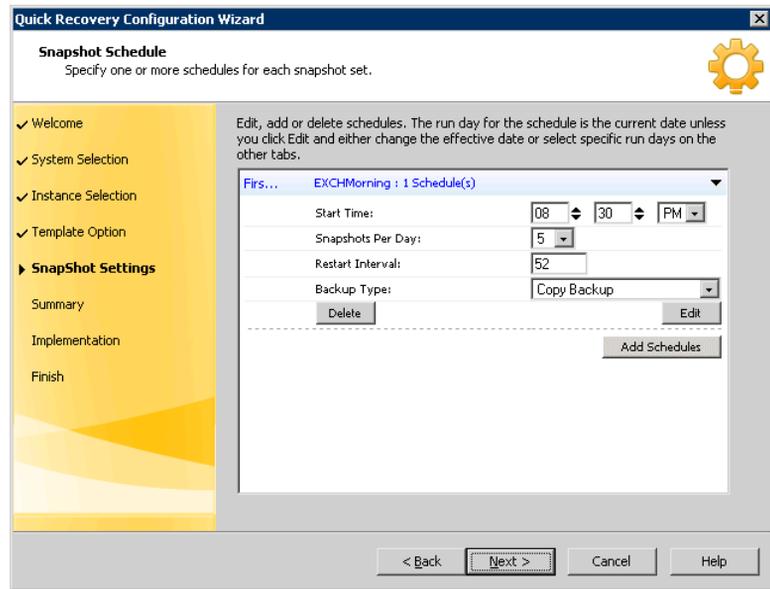
Mirror Preparation Time Click **Edit** and in the Mirror Preparation dialog box, specify the date and time for mirror preparation and click **OK**.

Preparing the snapshot mirrors slows down performance, so you may want to choose a time of low usage.

Snapshot Disks Assign one or more disks to each snapshot volume. Click the icon next to the disk column to select from available disks and to select a concatenated or striped volume layout.

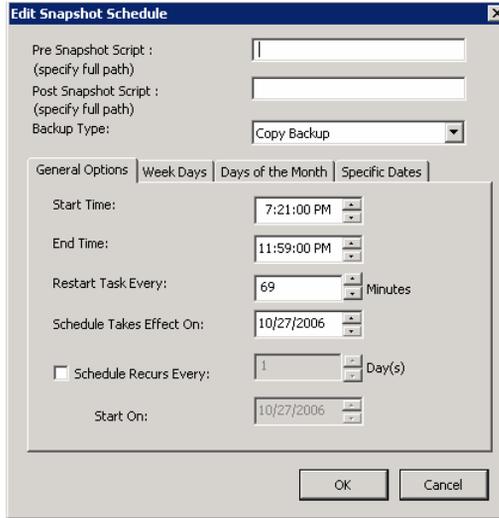
Mount	<p>Optionally, click Set to set a drive letter or mount path.</p> <p>The drive letter specified may not be available when the snapshot operation is performed. When this occurs, the snapshot operation is performed, but no drive letter is assigned.</p>
File path for snapshot XML file	<p>File path for the snapshot XML files. The path specified applies to all the snapshot sets that you are currently creating, for all currently selected databases. If the field is not editable, you are modifying an existing snapshot set schedule, so the value cannot be changed.</p> <p>For a non-clustered environment, the default location on Windows Server 2003 is:</p> <pre>C:\Documents and Settings\All Users\ Application Data\VERITAS\VSSXML\ application name</pre> <p>and on Windows Server 2008:</p> <pre>SystemDrive:\ProgramData\VERITAS\VSSXML\ application name</pre> <p>If a <code>redirect.txt</code> file has been created, the path in that file is shown in the field.</p> <p>For new schedules for a clustered server environment, if there is no <code>redirect.txt</code> file, the field is empty. Enter the full path to a location on shared storage, for example: <code>G:\SnapshotSets</code>.</p> <p>For Exchange 2010 configured in a VCS cluster, 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 <code>redirect.txt</code> file.</p> <p>You must store the XML files on a separate volume from the volumes that are included in snapshots.</p> <p>Note: If the snapshot XML files are not stored separately from the volumes that are included in the snapshots, a restore will fail.</p> <p>Ensure that you use a consistent location. That way, you can easily find the XML metadata file, which is required for recovery.</p>

- 10 Repeat [step 9](#) for any additional snapshot sets and when you finish, click **Next**.
- 11 In the Snapshot Schedule panel, choose one of the following:



- To specify a simple schedule, edit the default schedule settings shown on this panel: the time to start the snapshot (at least 1 hour later than the scheduled mirror preparation time), the number of snapshots to take, the interval in minutes between snapshots, and the type of snapshot. By default, the simple schedule takes effect on the current date and does not recur. Continue with [step 16](#).
- To specify a recurring schedule, a time window, a different run date, or other schedule details, click **Edit**.
Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.

- 12 If you clicked **Edit**, in the Edit Snapshot Schedule dialog box, choose from the following settings and then make additional selections on the dialog box tabs:



Pre Snapshot Script

Optionally, specify the full path of a script to run before the scheduled snapshot occurs.

In a cluster environment, script files should be located on shared storage.

For security purposes, ensure that script files are not accessible to any user with fewer privileges than the user account under whose context the scheduler service is running. Otherwise, a user with fewer privileges might change or replace the file that is launched by the scheduler service.

Post Snapshot Script

Optionally, specify the full path of a script to run after the snapshot is complete.

Backup Type

Select a backup type:

Full Backup is typically used for backup to tape or other storage media. It does the following:

- Creates a copy of the Exchange component
- Runs Eseutil to check for consistency before truncating the logs
- Truncates the transaction logs

Copy Backup is typically used for Quick Recovery. It creates a copy of the Exchange component, but does not truncate the transaction logs.

Copy Backup + Eseutil is used to create a copy and check the snapshot for consistency.

13 In the Edit Snapshot Schedule dialog box, on the General Options tab, you can specify the following:

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.

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 P.M. and an end time of 7 P.M., the snapshot occurs twice. If no interval is specified the snapshot occurs once.

Schedule takes effect on

The date on which the specified schedule takes effect. If you specify no other run day information on the other tabs, this date is used as the only run date.

If you specify run days on the other tabs, the schedule starts running on the earliest day specified. Therefore, if you want the current date included, specify it as an Include date on the Specific Dates tab.

Schedule recurs every	Enable this option to have the snapshot schedule continue to recur. Otherwise the schedule applies only for one day. Specify the number of days before restarting the snapshot schedule. For example, 1 day would mean the schedule takes effect daily, 2 days would mean every other day.
Start On	If you specify the Every option, specify the starting date.

- 14 In the Edit Snapshot Schedule dialog box, optionally make selections on the Week Days, Days of Month, and Specific Dates tabs as follows:

Week Days	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 1st 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 specify the last day of the 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 for the same snapshot set overlap for the same snapshot, 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.

- 15 When you are done making selections on the Edit Snapshot Schedule dialog box for this schedule, click **OK**.
- 16 In the Snapshot Schedule panel, choose from the following and when scheduling is complete for all snapshot sets, click **Next**:
- Edit schedules for any remaining snapshot sets
 - Click **Add Schedules** if you want to add a new schedule

- Click **Delete** if you want to remove a schedule
- 17 In the Summary panel, choose one or both of the following:
 - If you want to save the settings to a template file for reuse, click **Save** and save the template to a file location of your choice.
If you do not save the settings to a file, you can still view or modify them by launching the wizard and selecting the same instance and Exchange component.
 - If you are ready to implement the template with its current settings, click **Apply**.
If you click **Apply** without saving the template, you are prompted to confirm.
If you have saved the settings to a template file and want to exit the wizard without applying the template, click **Cancel**.
 - 18 In the Template Implementation panel, wait until the wizard shows that Quick Recovery configuration is complete and click **Next**.
 - 19 Click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt 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.
If you have just configured the schedules and they are not yet displayed, right-click the Storage Agent node and click **Refresh** to update the display.
- 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 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.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.

Situation

Resolution

You want to use the VSS Scheduler wizard but the Applications node is not shown in the VEA.

You may need to refresh the VEA display to see the node.

You are unable to locate a snapshot set XML file.

The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Drive letters assignments for snapshot volumes do not occur.

If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.

A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.

When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails.

To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.

Situation	Resolution
<p>Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.</p> <p>The following problems can occur:</p> <ul style="list-style-type: none">■ When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.■ Mirror preparation for the second database fails.	<p>You can choose from the following ways to avoid these problems:</p> <ul style="list-style-type: none">■ Include both databases in the snapshot when running the wizard the first time.■ Select a different disk for the second snapshot mirror when running the wizard the second time.■ Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.

In addition, when troubleshooting, you may want to review the following logs.

If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application  
Data\VERITAS\winsolutions\log
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting or modifying schedules

You can delete or modify a schedule that you created with the Quick Recovery Configuration Wizard by running the wizard again and deleting or editing the schedule on the Snapshot Schedule panel.

You cannot use the Quick Recovery Configuration Wizard to delete or modify schedules created from the VEA console.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings 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 **StorageAgent** node. Select and click to expand the **Scheduled Tasks > Exchange** node.
The scheduled snapshots are listed on the right pane.
- 3 Right-click the name of the snapshot schedule and select **Delete Schedule**.

Synchronizing schedules after adding a cluster node

In a cluster environment, you may add a cluster node after you set up snapshot schedules with the Quick Recovery Configuration Wizard.

In such a case, you can ensure that the schedules are available on the new node by running the Quick Recovery Configuration Wizard again.

To synchronize schedules after adding a node

- 1 Start the Quick Recovery Configuration Wizard from the Solutions Configuration Center.
- 2 Continue through the wizard until the Synchronizing Schedules panel shows that synchronization between cluster nodes is complete.
- 3 Click **Cancel** to exit the wizard.

Scheduling or creating an individual snapshot set

This chapter includes the following topics:

- [About scheduling or creating an individual snapshot set](#)
- [Tasks to schedule a new snapshot](#)
- [Tasks to create a one-time snapshot set](#)
- [Reviewing the prerequisites](#)
- [Preparing the snapshot mirrors](#)
- [Scheduling a snapshot set](#)
- [Troubleshooting scheduled snapshots](#)
- [Creating a one-time snapshot set](#)
- [Refreshing a snapshot set manually](#)

About scheduling or creating an individual snapshot set

Typically you set up your initial snapshot schedules with the Quick Recovery Configuration Wizard.

See [Chapter 6, “Implementing snapshot sets with the configuration wizard”](#) on page 69.

However, later, you may want to schedule an additional snapshot. For example, you configured a daily snapshot for Quick Recovery use and now want to add a weekly snapshot on a different disk for backup use.

You can use the VSS Exchange Snapshot Scheduler Wizard to add a snapshot schedule for a selected storage group. Like the Quick Recovery Configuration

Wizard, the scheduler wizard enables you to automate the refreshing of snapshots according to the schedule that you define.

However, unlike the Quick Recovery Configuration Wizard, the VSS Exchange Snapshot Scheduler Wizard does not prepare snapshot mirrors. You must use the Prepare command to prepare the snapshot mirrors before running the VSS Exchange Snapshot Scheduler Wizard for that storage group. In addition, you can use the scheduler wizard to schedule only one snapshot set for one storage group at a time.

Note: Adding a snapshot set using the VSS Exchange Snapshot Scheduler will not update the template settings created with the Quick Recovery Configuration Wizard. If you want to keep the template settings up to date, you should instead run the Quick Recovery Configuration Wizard to modify the schedule.

See “[Scheduling a snapshot set](#)” on page 93.

At times you may want to create a one-time snapshot set. You can do so using either the vxsnap command line utility or from the VEA console using the VSS Exchange Snapshot and Snapback wizards.

See “[Creating a one-time snapshot set](#)” on page 104.

Note: To snapshot only a single volume rather than multiple volumes, you can use the VEA Snapshot Volume Wizard. See the *Veritas Storage Foundation Administrator’s Guide*.

Tasks to schedule a new snapshot

[Table 7-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 7-1 Tasks for scheduling a new snapshot set

Objective	Tasks
“ Reviewing the prerequisites ” on page 91	■ Verifying hardware and software prerequisites
“ Preparing the snapshot mirrors ” on page 92	■ Creating snapshot mirrors using the using the VEA Prepare command or the vxsnap utility
“ Scheduling a snapshot set ” on page 93	■ Using the VSS Exchange Snapshot Scheduler Wizard to create the initial snapshot and set up the schedule for keeping it refreshed.

Tasks to create a one-time snapshot set

[Table 7-2](#) outlines the high-level objectives and the tasks to complete each objective.

Table 7-2 Tasks for creating a one-time snapshot set

Objective	Tasks
“Reviewing the prerequisites” on page 91	■ Verifying hardware and software prerequisites
“Preparing the snapshot mirrors” on page 92	■ Creating snapshot mirrors using the VEA Prepare command or the vxsnap utility
“Creating a one-time snapshot set” on page 104	■ Creating the one-time snapshot set using the VEA or the vxsnap utility.
“Refreshing a snapshot set manually” on page 107	■ Refreshing the one-time snapshot set if necessary.

Reviewing the prerequisites

Before implementing an individual snapshot set, you should do the following:

- Ensure that your system hardware and software meets the requirements. See [“System requirements”](#) on page 37 in [Chapter 4, “Planning a Quick Recovery snapshot solution”](#).
- Set up your Exchange storage groups for use with Storage Foundation for Windows. See [Chapter 5, “Configuring Exchange for Quick Recovery snapshots”](#) on page 53.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- The vxsnap commands must be invoked on a local system. On Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Preparing the snapshot mirrors

To prepare the snapshot mirrors in order to create a snapshot from the VEA or from the `vxsnap` command line, you can use either of the following methods:

- The Prepare command from the VEA
You repeat the VEA console Prepare operation for each volume in the storage group.
- The `vxsnap prepare` command from the CLI
Use the `vxsnap prepare` command to prepare a mirror for each of the volumes associated with the databases or transaction logs of your Exchange storage group.

The snapshot mirrors remain attached to the original volumes and continue to be updated until you use the VSS Exchange Snapshot Wizard, the `vxsnap create` command, or the VSS Exchange Snapshot Scheduler Wizard to create the snapshot set.

For the snapshot volumes, make sure to select disks or LUNs that are not used for production data. You can create more than one snapshot volume on the same disk as long as there is sufficient space available and as long as the snapshots are of different production volumes.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare** option from the context menu.
- 2 Choose one of the following.

If the volume is not mirrored

Choose **Manually select disks**, use the **Add** and **Remove** buttons to move the desired disks to the Selected disks box, and click **OK**.

If the volume is mirrored and no additional disks are available to create a new mirror

Click on an existing plex and click **OK**.

If the volume is mirrored and there are additional disks available on your system

Choose either to use an existing mirror for the snapshot or to create a new mirror.

- To create a new mirror, click **Select Disk**, use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box, and click **OK**.
- To use an existing mirror, click **Select existing mirror for snap**, select the desired mirror, and click **OK**.

- 3 Repeat [step 1](#) and [step 2](#) to create a snapshot mirror for each volume associated with the Exchange storage group and transaction logs. Verify that the lower pane of the VEA console indicates that the resynchronization process is complete before creating the snapshot set.

To create the snapshot mirrors using the command line

Type the command, as in the following example:

```
vxsnap prepare component=SG2/writer="Microsoft Exchange Writer"  
source=G:/harddisk=harddisk4  
source=H:/harddisk=harddisk4  
source=I:/harddisk=harddisk5
```

In this example, snapshot mirrors for the database volumes mounted at G: and H: are created on disk 4 and a snapshot mirror for the log volume mounted at I: is created on disk 5.

Note: For Exchange 2007, if local continuous replication (LCR) is enabled and you want to create a snapshot mirror for the passive copy, specify the writer option as "Microsoft Exchange Writer Replica."

The complete syntax of the `vxsnap prepare` command is:

```
vxsnap prepare component=<componentName>/writer=<writerName>  
[-b] [source=<volume>/harddisk=<harddisk,...>] ...]
```

See "[vxsnap prepare](#)" on page 171.

Scheduling a snapshot set

Before you run the VSS Exchange Snapshot Scheduler Wizard to schedule a snapshot set for a storage group, you must prepare snapshot mirrors attached to the volumes in the storage group.

See "[Preparing the snapshot mirrors](#)" on page 92.

You can then use the VSS Exchange Snapshot Scheduler Wizard to schedule the initial snapshot and to set up the schedule for keeping it refreshed.

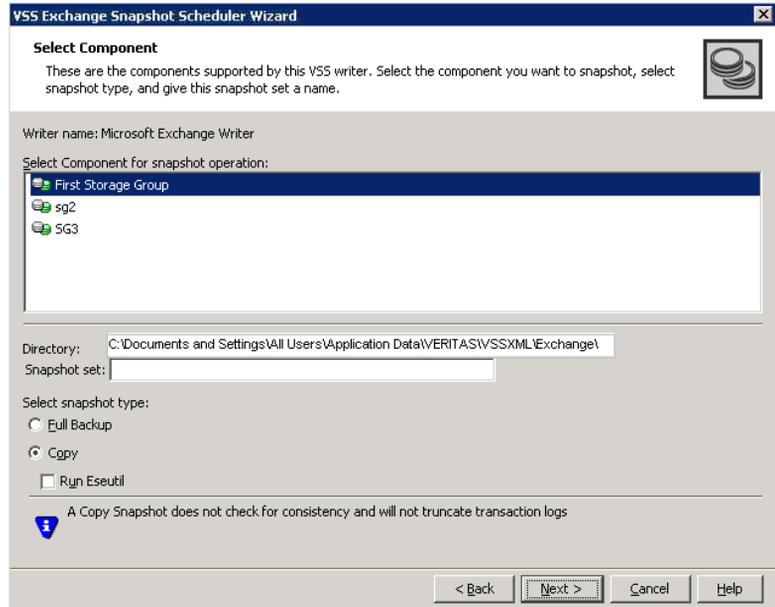
When the scheduled snapshot occurs, the snapshot mirrors are detached from the original volumes, creating separate on-host snapshot volumes as well as an XML file to store the Exchange and snapshot volume metadata. The scheduled process integrates with VSS to quiesce the databases and then simultaneously snapshot the volumes in the storage group. This snapshot is done while the databases are online and without disrupting the email flow.

Note: For Exchange 2007, if using local continuous replication (LCR), you can schedule snapshots only for the active copy, not the passive copy.

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 Right-click **Exchange** 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 2003 or 2007, the component is a storage
group.

Directory

Accept or enter a directory location for the XML file.
The XML file is stored by default in the directory
shown on the screen.

Note: The XML file for the snapshot must be stored
separately from the volumes that are included in the
snapshots, otherwise a restore will fail.

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
`redirect.txt`. This text file should contain a single
text line specifying the full path to the location of the
XML file, for example, `G:\BackupSets`. Save the
`redirect.txt` file in the default directory
`C:\Program Files\Veritas\Veritas Volume
Manager 5.1\VSSXML` on each node of the cluster.

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 will fail.

Snapshot set

Enter a name for the snapshot set.

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.

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 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 drive letters specified may not be available when the snapshot is taken. When this occurs, the snapshot operation is performed, but no drive letters are assigned.

Plex

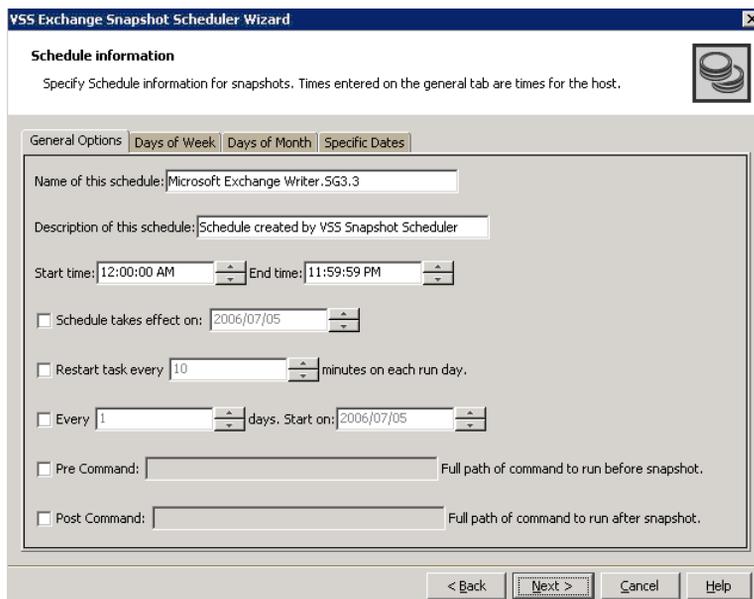
Optionally, click a plex and select a new choice from the drop-down menu.

7 Optionally, in the Synchronized Snapshot panel (VVR 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 Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

See *Veritas 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, the component 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 P.M. and an end time of 7 P.M, the snapshot occurs twice.</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 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> <p>Note: Precommands which launch windows or any other GUI related entities are not allowed in the Scheduler.</p>
Post Command	<p>Optionally, specify the full path of a command script to run after the snapshot is complete.</p> <p>Note: Postcommands which launch windows or any other GUI related entities are not allowed in the 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 checkbox to schedule the snapshot for the last day of each month.</p>

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.

10 Click **Next**.

11 Review the snapshot set and schedule details and click **Finish**.

Viewing the status of scheduled snapshots

If a scheduled snapshot fails for some reason, the scheduler process will attempt 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.
If you have just configured the schedules and they are not yet displayed, right-click the Storage Agent node and click **Refresh** to update the display.
- 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 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.

Troubleshooting scheduled snapshots

When scheduling snapshots using the Quick Recovery Configuration Wizard or the VSS Snapshot Scheduler Wizard, you can use the information in the following table to help avoid problems or troubleshoot situations.

Situation	Resolution
Snapshots do not occur on a scheduled date.	The date may have been excluded under “Specific Dates” on the schedule. Excluding a date takes precedence over other days scheduled.
A snapshot is not taken on the date that you create the schedule, although that date is specified as the “Schedule takes effect on date”.	The date shown in the Schedule takes effect on field is used as a run day only if no other run days are specified. If you specify other run days, the schedule starts running on the earliest day specified. If you want the current date included in a schedule, specify it as an Include date on the Specific Dates tab.
A scheduled snapshot does not occur for an imported template schedule.	If you import a template that uses the current date as the “takes effect on” date, but part of the schedule cannot be implemented for the current date, the effective date is shifted to the following day. For example, if the schedule includes two daily snapshots, one in the morning and one at night, and you import the schedule in the afternoon, neither of the snapshots occur on the current date. Both occur the following day.
While running the Quick Recovery Wizard to modify template settings, the existing settings are not displayed.	If a scheduled snapshot operation is occurring on the selected object at the same time that you select the object in the wizard, the wizard is temporarily unable to display the existing settings. You can try running the wizard later.
A schedule is unavailable to be edited in the Quick Recovery Wizard.	Expired schedules cannot be edited. Instead, delete the expired schedule and add a new one.
You want to use the VSS Scheduler wizard but the Applications node is not shown in the VEA.	You may need to refresh the VEA display to see the node.
You are unable to locate a snapshot set XML file.	The VSS Snapshot Scheduler Wizard assigns a prefix of “VM_” to the name you assign.

Situation	Resolution
Drive letters assignments for snapshot volumes do not occur.	If time elapses between when you use the wizard to assign drive letters and when the snapshot operation occurs, a drive letter you assigned may become unavailable. When this occurs, the snapshot operation is performed, but no drive letters are assigned. You can assign or change drive letters or mount paths in Veritas Enterprise Administrator.
A scheduled snapshot fails after you have done a manual snapback (reattach) of one or more of the snapshot volumes in the snapshot set.	When a snapback is done manually for a volume rather than by using the VSS Snapback wizard or allowing the scheduler to handle it, the XML metadata file is not deleted. To take the next scheduled snapshot, the scheduler attempts to reattach all the volumes specified in the XML metadata file and fails because one or more are already reattached. Therefore, the snapshot also fails. To ensure that the next scheduled snapshot works correctly, use the VSS Snapback wizard (or the vxsnap utility) to reattach using the XML file.
Two databases are on the same volume. You run the Quick Recovery wizard to schedule the first database. You then run it again to schedule a snapshot for the second database on the same disk. But the mirror preparation for the first database is not yet complete.	You can choose from the following ways to avoid these problems: <ul style="list-style-type: none">■ Include both databases in the snapshot when running the wizard the first time.■ Select a different disk for the second snapshot mirror when running the wizard the second time.■ Wait for the mirror preparation for the first snapshot to finish before running the wizard to schedule the second snapshot.
The following problems can occur: <ul style="list-style-type: none">■ When running the wizard, the second time, it does not update the available disk space shown for the disk where the first snapshot is scheduled.■ Mirror preparation for the second database fails.	

In addition, when troubleshooting, you may want to review the following logs. If a schedule fails, check the scheduler service logs in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager  
5.1\logs\SchedService.log
```

If a snapshot fails, check the VxSnap.log file in the following folder:

```
C:\Program Files\Veritas\Veritas Volume Manager 5.1\logs
```

Quick Recovery Configuration wizard log files are located in the following paths.

For Windows Server 2003:

```
C:\Documents and Settings\All Users\Application  
Data\VERITAS\winsolutions\log
```

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```

Deleting schedules

You can delete (but not modify) a schedule from the VEA console. To modify a schedule, run the wizard again and select the same instance and component.

Note: Deleting a snapshot schedule using the VEA does not update template settings created with the Quick Recovery Configuration Wizard.

Note: You cannot modify a schedule that has expired.

You can also delete (but not modify) a schedule from the VEA console.

Note: The VEA can delete snapshot schedules only; it does not delete mirror preparation scheduled with the Quick Recovery Configuration Wizard. In addition, deleting a snapshot schedule using the VEA does not update template settings 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 **StorageAgent** node. Select and click to expand the **Scheduled Tasks > Exchange** node.
The scheduled snapshots are listed on the right pane.
- 3 Right-click the name of the snapshot schedule and select **Delete Schedule**.

Creating a one-time snapshot set

Creating a one-time snapshot set is a two-step process:

- The first step is to prepare snapshot mirrors attached to all the original volumes in the specified storage group. If you are creating a snapshot set after a snapback to refresh existing snapshot mirrors, you can skip this step. See [“Preparing the snapshot mirrors”](#) on page 92.
- The second step uses either the `vxsnap create` command or the VSS Exchange Snapshot Wizard to create the snapshot set by detaching the snapshot mirrors from the original volumes. This step creates separate on-host snapshot volumes as well as an XML file to store the Exchange and snapshot volume metadata.

The `vxsnap create` command and VSS Exchange Snapshot Wizard integrate 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. The resulting snapshot set provides a complete picture of the storage group at the point in time the command is issued.

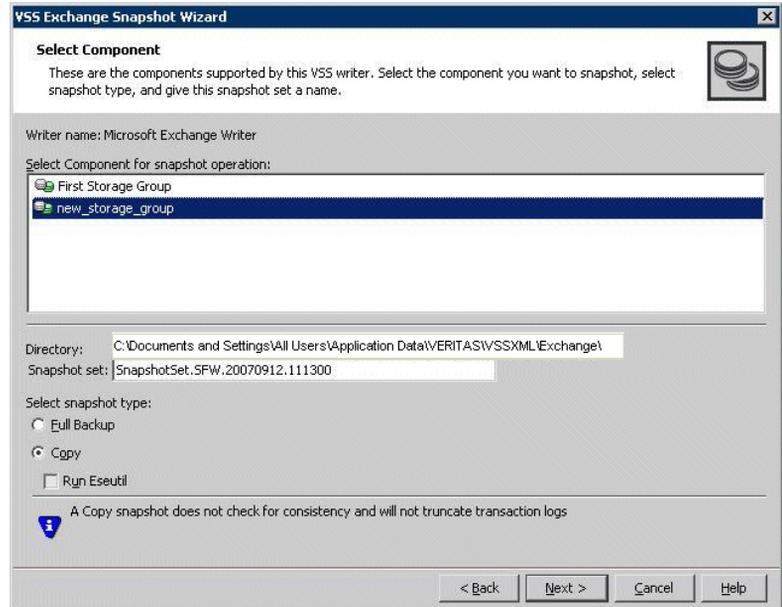
The VSS Exchange Snapshot Wizard can be run from either a local system or a remote node. The `vxsnap` utility must be run from the local system.

See [“Scheduling a snapshot set”](#) on page 93.

To create the snapshot set from the VEA console

- 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 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange 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**.
- 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 2003 or 2007, the component is a storage
group.

Directory	<p>Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory 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 will fail.</p> <p>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 "redirect.txt." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, G:\BackupSets. Save the redirect.txt file in the default directory C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML on each node of the cluster.</p> <p>Note: 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 will fail.</p>
Snapshot set	<p>Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.</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	<p>In the Change Attributes panel, optionally change the attributes for the snapshot volumes as follows and click Next:</p>
Snapshot Volume Label	<p>Displays the read-only label for the snapshot volume.</p>

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 Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

To create the snapshot set from the command line

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]  
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<newVolName>] [/Plex=<plexName>] ...writer=<writerName>  
component=<componentName> [backuptype=<backuptype>] [-E] [-O] [-C] [  
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Creating the snapshot set with the copy backup type does not automatically run the Eseutil consistency check nor truncate the transaction logs. You can check the snapshot for consistency by specifying the `-E` flag.

Note: Any text string that contains spaces must be enclosed in quotation marks.

See “[vxsnap create](#)” on page 172.

Refreshing a snapshot set manually

Once a snapshot set has been created, it can be refreshed quickly since the time-consuming step of preparing the mirrors is not required.

Normally, if you want to periodically refresh a snapshot set, you set up the snapshot schedule using the VSS Exchange Snapshot Scheduler Wizard or the Quick Recovery Configuration Wizard.

However, if you should need to manually refresh a snapshot set, you can do so. To refresh the snapshot set requires the following tasks:

- “[Reattaching the split-mirror snapshots](#)” on page 108
- “[Creating the refreshed snapshot set](#)” on page 110

Note: The VSS Refresh option available in the VEA console from the Microsoft Exchange Writer object refreshes the display of the VSS Writer and components. It does not refresh the snapshot set.

Reattaching the split-mirror snapshots

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot set so that it matches the current state of its original Exchange mailbox storage group. The wizard is available in the context menu of the VSS Writer object.

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 be accessing the snapshot set.
- 2 From the VEA console URL bar, select the *<host name>* which is 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 on the node of the 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 required information needed 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**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command, as in the following example:

```
vxsnap -x Image1.xml reattach writer="Microsoft Exchange Writer"
```

The complete syntax for the `vxsnap reattach` command is:

```
vxsnap -x <filename> [-f][-b] reattach writer=<writername>  
[secHosts=<secondary hosts>]
```

See [“vxsnap reattach”](#) on page 175.

Creating the refreshed snapshot set

Create a new snapshot set of the storage group using either the VSS Exchange Snapshot Wizard or the `vxsnap create` command.

See “[Creating a one-time snapshot set](#)” on page 104.

Recovering Exchange mailbox storage group or databases

This chapter includes the following topics:

- [About recovery using Quick Recovery snapshots](#)
- [Tasks for recovery using Quick Recovery snapshots](#)
- [Prerequisites for recovery](#)
- [Recovery using an Exchange 2007 passive copy snapshot](#)
- [Recovering using the VSS Restore Wizard](#)
- [Recovering using the vxsnap utility](#)
- [Post-recovery steps](#)

About recovery using Quick Recovery snapshots

The on-host snapshot set can be used to quickly recover a storage group after logical corruption. You can restore a storage group either to the Point in Time (PIT) that the `vxsnap restore` command was last refreshed or, using current transaction logs, to the Point of Failure (POF) of the database. Additionally, you can recover a single database to the point of failure.

You can use either the VSS Restore Wizard from the VEA console or the `vxsnap restore` command. Both the VSS Restore wizard and `vxsnap restore` command integrate with VSS to notify the Exchange VSS Writer to prepare for the restore before the snapback operation and then to complete post-restore processes afterwards.

After completing the recovery, you refresh the snapshot set.

Tasks for recovery using Quick Recovery snapshots

[Table 8-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 8-1 Tasks for recovery using Quick Recovery snapshots

Objective	Tasks
“Prerequisites for recovery” on page 113	■ Verifying hardware and software prerequisites
“Recovery using an Exchange 2007 passive copy snapshot” on page 113	■ (Exchange 2007 only) Reviewing prerequisites for using a passive copy snapshot for recovery
“Recovering using the VSS Restore Wizard” on page 115	■ Running the VSS Restore Wizard from the VEA console to restore a storage group to point in time or restore a storage group or database point of failure ■
“Recovering using the vxsnap utility” on page 120	■ Using the vxsnap utility to restore a storage group to point in time or restore a storage group or database point of failure ■
“Post-recovery steps” on page 123	■ Refreshing the snapshot set ■ Performing additional tasks in a VVR environment

Prerequisites for recovery

You can use the VSS Restore Wizard or vxsnap utility to recover a storage group or single database in cases where a snapshot set is available and where there has been no hardware failure.

Caution: Do not attempt to use the VSS Restore Wizard for recovery after hardware failure. Instead, refer to the special procedures for recovery after hardware failure.

See [Chapter 9, “Recovering after hardware failure”](#) on page 125.

For Exchange 2007, see the following additional information before you begin:

- Before restoring an Exchange 2007 database in a Microsoft clustering environment (MSCS or Failover Clustering):
 - Manually dismount the stores and set them for overwrite by restore.
 - If CCR is enabled, disable the circular logging feature before restoring the database.
 - If CCR is enabled, suspend CCR before restoring the database.
- Before restoring from a passive copy snapshot, additional steps are required. See [“Recovery using an Exchange 2007 passive copy snapshot”](#) on page 113.
- If you have set up a recovery storage group (RSG) in Exchange 2007, you have the option to restore to the RSG rather than to the production volume, leaving the production volume intact. The RSG must be created before you begin the recovery. Follow the Microsoft instructions in creating the RSG. The database names in the recovery storage group must match those in the source storage group.

Recovery using an Exchange 2007 passive copy snapshot

For Exchange 2007, if you used SFW to create a snapshot of a passive copy of the storage group, the passive copy snapshot can be used to restore an Exchange database.

See the following topics:

- [“Overview of recovery steps using an Exchange 2007 passive copy snapshot”](#) on page 114
- [“Example of restoring from a passive copy”](#) on page 115

Overview of recovery steps using an Exchange 2007 passive copy snapshot

To recover using an Exchange 2007 passive copy snapshot, you follow this process:

- Use the Microsoft Exchange Management Shell to execute the following cmdlets:
 - Dismount-Database
 - Restore-StorageGroupCopy
The LCR copy is automatically disabled when running the Restore-StorageGroupCopy cmdlet. You can enable it again once all recovery steps are complete.
 - Mount-Database
Refer to the Microsoft Exchange documentation for additional information on Microsoft Exchange Management Shell and cmdlets.
For an example of how to use these cmdlets, see the following topic:
 - [“Example of restoring from a passive copy”](#) on page 115
- Execute the Refresh command from the VEA or command line.
- In a Microsoft clustering environment (MSCS or Failover Clustering), ensure that you dismount the stores and set them for overwrite by restore. In a VCS or SFW environments, SFW attempts to dismount the stores before beginning the restore operation. If SFW fails to dismount the stores, dismount them manually.
- If a snapshot schedule exists for the active storage group, it is now invalid, and you must delete it before beginning the restore operation.
- Use either the VEA wizard or the vxsnap CLI utility with the active writer (Microsoft Exchange Writer) to restore the passive copy snapshot. See the following topics:
 - [“Recovering using the VSS Restore Wizard”](#) on page 115.
 - [“Recovering using the vxsnap utility”](#) on page 120

Note: You can restore using either Point in Time or Point of Failure. Restoring just the database log files from a passive copy is not supported.

Example of restoring from a passive copy

The following example illustrates a Point in Time recovery from a snapshot of a passive copy of an Exchange storage group, SG1, that contains two databases, DB1 and DB2, on an Exchange server, TestExch.

- 1 From the Exchange Management Shell, run the Dismount-Database cmdlet on the DB1 and DB2 databases.

```
Dismount-database -Identity TestExch\SG1\DB1  
Dismount-database -Identity TestExch\SG1\DB2
```

- 2 Run the Restore-StorageGroupCopy cmdlet on the SG1 storage group.

```
Restore-StorageGroupCopy -Identity TestExch\SG1  
-ReplaceLocations
```

- 3 Run the Mount-Database cmdlet on the DB1 and DB2 databases.

```
Mount-database -Identity TestExch\SG1\DB1  
Mount-database -Identity TestExch\SG1\DB2
```

- 4 Perform a vxsnap refresh.

```
vxsnap refresh
```

- 5 Perform the VSS restore operation using the snapshot of the passive copy.

```
vxsnap -x snapdata.xml restore RestoreType=PIT  
writer="Microsoft Exchange Writer"
```

Note: For this example, assume that the snapshot of the passive copy was performed with

```
vxsnap -x snapdata.xml create writer="Microsoft  
Exchange Writer Replica" component=SG1 backupType=COPY  
-E -O
```

Recovering using the VSS Restore Wizard

Using the VSS Restore Wizard, you can do the following:

- Restore a storage group to the point in time of the snapshot set
- Perform a roll-forward recovery of a storage group or single database to the point of failure

Note: Recovery Database for Exchange 2010 is not supported by SFW.

Before you begin, review the prerequisites.

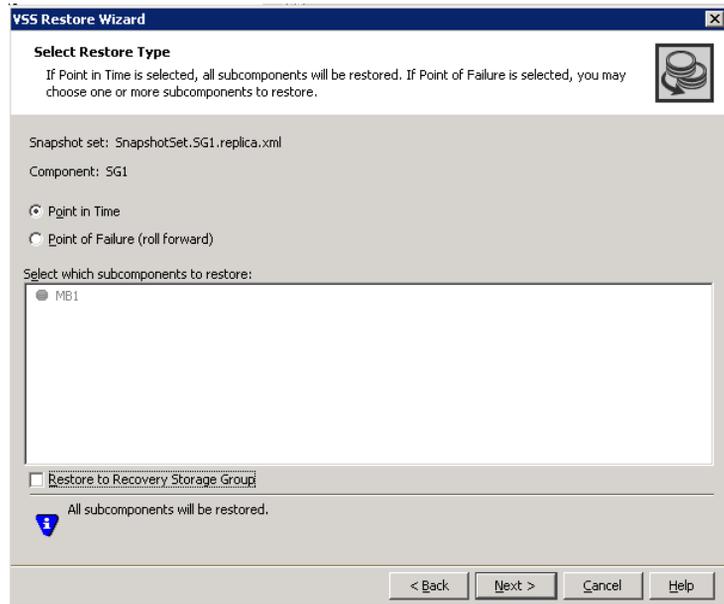
See “[Prerequisites for recovery](#)” on page 113.

Note: SFW dismounts the stores before beginning the restore operation. If it fails to dismount the stores, the restore operation fails. In that case, manually dismount the stores and set them for overwrite by restore. If LCR is enabled, suspend LCR. Then repeat the restore operation. If you suspended LCR, re-enable it after the restore operation.

To restore to the point in time of the 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 be accessing 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 wish to restore and click **Next**.
The XML metadata file contains all required information 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 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 Exchange Snapshot Wizard or `vxsnap create` command to create a new snapshot set of all the volumes in the storage group.

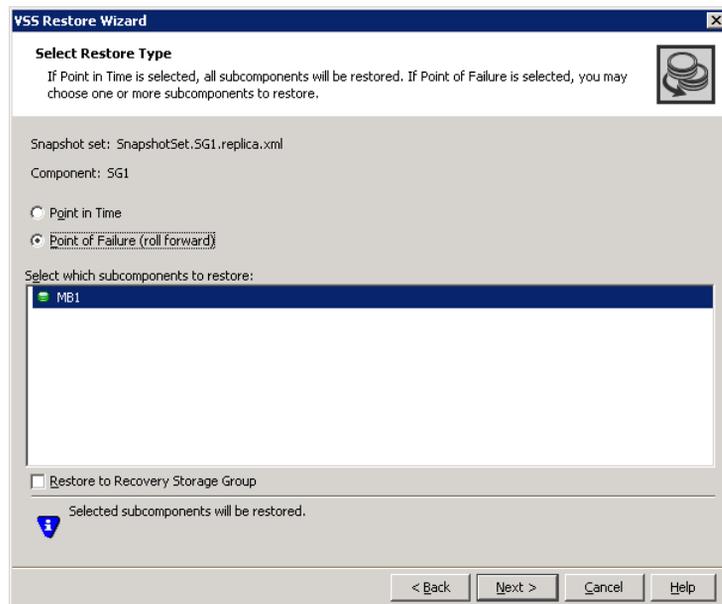
To perform a roll-forward recovery of one or more databases to the point of failure

- 1 Dismount all the databases (stores) in the Exchange storage group:
 - For Exchange 2003, you can use the Exchange System Manager.
 - For Exchange 2007, you can use the Exchange Management Console.
- 1 Close the database application GUI and all Explorer windows, applications, consoles (except the VEA console), or third-party system management tools that may be accessing 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 wish 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**

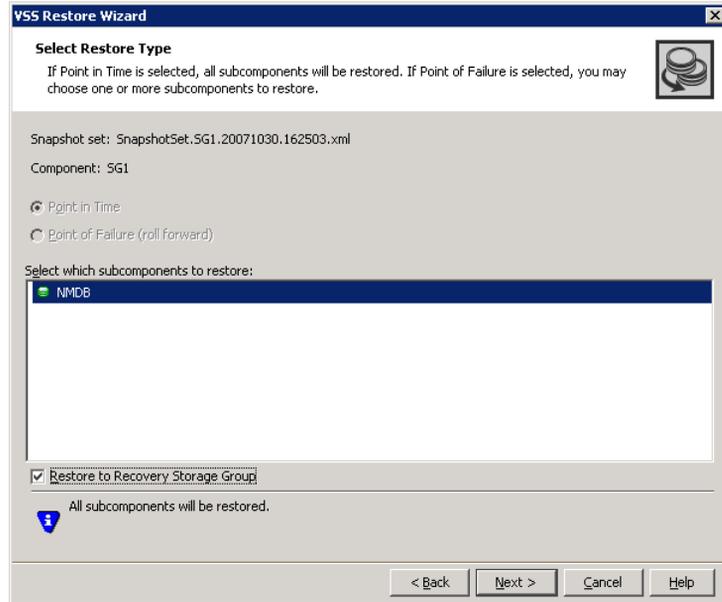
- 8 Verify that the restore specifications are correct and click **Finish**.
- 9 Mount all databases (stores) in the Exchange storage group; the transaction logs will roll forward to the point of failure. This action can be observed in the Event Viewer Application log.

- 10 To refresh the snapshot set at this point, 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 be accessing 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 wish to restore and click **Next**.
The XML metadata file contains all required information 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**

- 8 Verify that the restore specifications are correct and click **Finish**.
- 9 Mount all the databases (stores) in the Exchange storage group.

Recovering using the vxsnap utility

You can also recover exchange mailbox databases by using the vxsnap utility from the Command Line Interface (CLI).

Refer to Microsoft Exchange Shell cmdlets for details.

Note: SFW automatically dismounts the stores before beginning the restore operation. If it fails to dismount the stores, the restore operation fails. In that case, manually dismount the stores and repeat the restore operation. For a VCS cluster setup, it offlines and onlines the resource for the database automatically. If automatic offlining of the resource fails, then manually offline the resource.

Using the vxsnap utility, you can do the following:

- Restore a storage group to the point in time (PIT) of the snapshot set

- Perform a roll-forward recovery of a storage group to the point of failure (POF)
- Perform a roll-forward recovery of a single database to the point of failure
- Exchange 2007 only: Restore to a recovery storage group (RSG) rather than the production store

Before you begin, review the prerequisites.

See “[Prerequisites for recovery](#)” on page 113.

You use the `vxsnap restore` command. The complete syntax of the `vxsnap restore` command is:

```
vxsnap -x filename [-f] [-b] [-r] restore restoreType=<PIT|POF>
[subComponent=subComponentName] [writer=WriterName]
[RSG=<YES|NO>]
vxsnap -x <filename>[-f] [-b] [-r] [-a] restore
restoreType=<PIT|POF>writer=<writename> [subComponent=<subComponentName>] [RSG=<Yes|No>]
```

To restore a storage group to the point in time of the snapshot set using the `vxsnap restore` command

- 1 Close the database GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Click **Start > Run**. In the command window, type the `vxsnap restore` command

```
vxsnap -x Image1.xml restore restoretype=PIT writer="Microsoft Exchange Writer"
```

The XML metadata file contains all the necessary information about the Exchange storage group, including the names of the database and transaction logs volumes.

- 3 Mount all the databases (stores) in the Exchange storage group.
- 4 To refresh the snapshot set use the `vxsnap create` command to create a new snapshot of all the volumes in the storage group.

To perform a roll-forward recovery of a storage group to the point of failure using the `vxsnap restore` command

- 1 Close the database GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Click **Start > Run**. In the command window, type the `vxsnap restore` command

```
vxsnap -x Image1.xml restore restoretype=POF writer="Microsoft Exchange Writer"
```

In this example, `Image1.xml` is the name of the XML metadata file associated with the snapshot set.

If you want to restore to an RSG (Exchange 2007 only), use the `RSG=Yes` option.

- 3 Mount all databases (stores) in the Exchange storage group; the transaction logs will roll forward to the point of failure. This action can be observed in the Event Viewer Application log.
- 4 To refresh the snapshot set at this point, 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 perform a roll-forward recovery of a single database to the point of failure using the `vxsnap restore` command

- 1 Close the database GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.

- 2 Type the `vxsnap restore` command, as in the following example:

```
> vxsnap -x Image1.xml restore restoretype=POF  
      subcomponent=DB1 writer="Microsoft Exchange Writer"
```

In this example, DB1 is the name of the Exchange database that is associated with the volume SG2-DB1, and Image1.xml is the name of the XML metadata file associated with the snapshot set.

If you want to restore to an RSG (Exchange 2007 only), use the `RSG=Yes` option.

- 3 Mount all databases (stores) in the Exchange storage group; the transaction logs will roll forward to the point of failure. This action can be observed in the Event Viewer Application log
- 4 Mount all databases; the transaction logs will roll forward to the point of failure. This action can be observed in the Event Viewer Application log.
- 5 To refresh the snapshot set at this point, 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.

Post-recovery steps

After you have performed any of the recovery methods, whether point-in-time or roll forward, you should refresh the snapshot set, by performing a snapback to reattach the snapshot mirrors, and then create a new snapshot set.

See “[Refreshing a snapshot set manually](#)” on page 107 in [Chapter 7, “Scheduling or creating an individual snapshot set”](#).

In a VVR environment, there is an additional post-recovery step. During a point-in-time recovery in a VVR environment, the volumes on the secondary site lose write-order fidelity. DCM automatically becomes active to ensure data consistency between the primary and secondary sites. While DCM is active, the volumes cannot be expanded by either manual or AutoGrow operations. You must perform a manual resynchronization of the secondary to deactivate DCM.

To resynchronize the node at the secondary site with the node at the primary site:

- 1 Right-click on the primary RVG and click **Resynchronize Secondaries**.
- 2 Click **Yes** to resynchronize the nodes.

Recovering after hardware failure

This chapter includes the following topics:

- [About recovery after hardware failure](#)
- [Tasks for recovering after hardware failure](#)
- [Reviewing the prerequisites](#)
- [Reviewing the sample configuration](#)
- [Scenario I: Database and transaction logs volumes are missing](#)
- [Scenario II: Database volumes missing, transaction logs are available](#)
- [Scenario III: Some DB volumes missing, transaction logs are available](#)
- [Refreshing the snapshot set](#)
- [Refreshing the snapshot set on the current disks](#)
- [Moving the production volumes to different disks and refreshing the snapshot set](#)

About recovery after hardware failure

A system failure that results in the loss of either database or transaction log volumes leads to unplanned downtime. If the failure does not affect the disk or disks where the snapshot set resides, you can recover from the failure and resume normal email operations faster by using the snapshot set than you could by restoring from your backup media.

You can use the `vxsnap` utility to recover a storage group after production volumes are lost due to a hardware failure.

Note: The methods described in this chapter are the simplest possible and assumes that a volume snapshot set is already available before proceeding with the restore operation. However, if the snapshot set is not available, you will need to recover it from your tape backup.

Recovery scenarios

Table 9-1 lists the storage group recovery scenarios that you choose from if the complete snapshot set (including the XML metadata file) is available:

Table 9-1 Available recovery type after missing production volume

Scenario	Database Volumes	Transaction Logs Volume	Recovery
Scenario I	One or more volumes are missing	Missing	point in time
Scenario II	All volumes are missing	Available	point in time or point of failure
Scenario III	One or more volumes are missing. At least one volume is available	Available	point of failure

Caution: If you choose to restore the storage group to the point in time, you cannot later restore it to the point of failure. You can only perform one recovery procedure on a mailbox database.

Recovery tasks

Perform the recovery tasks in the order shown below.

For a VSS-integrated recovery, you need the following information:

- Prepare for the recovery
 - Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
 - Use Exchange Management Console 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.
- Use the `vxsnap restore` command to recover the databases.
- Refresh the snapshot set.

Tasks for recovering after hardware failure

Table 9-2 outlines the high-level objectives and the tasks to complete each objective.

Table 9-2 Tasks for recovering after hardware failure

Objective	Tasks
“Reviewing the prerequisites” on page 128	<ul style="list-style-type: none"> ■ Reviewing the prerequisites for a recovery after system failure
“Scenario I: Database and transaction logs volumes are missing” on page 129	<ul style="list-style-type: none"> ■ Preparing for the recovery ■ Changing drive letters or mount points of all the snapshot volumes ■ Restoring the storage group to the point in time ■ Refreshing the snapshot set
“Scenario II: Database volumes missing, transaction logs are available” on page 137	<ul style="list-style-type: none"> ■ Preparing for the recovery ■ Removing the drive letter or mount point of the transaction logs volume ■ Changing drive letters or mount points of all the snapshot database volumes ■ Restoring the storage group to the point in time or Recovering the storage group to the point of failure ■ Refreshing the snapshot set
“Scenario III: Some DB volumes missing, transaction logs are available” on page 145	<ul style="list-style-type: none"> ■ Preparing for the recovery ■ Changing drive letters or mount points of snapshot database volumes for missing database volume only ■ Recovering the storage group to the point of failure ■ Refreshing the snapshot set
“Refreshing the snapshot set” on page 153	<ul style="list-style-type: none"> ■ Understanding your current configuration ■ Choosing the disks for the refreshed snapshot sets.

Table 9-2 Tasks for recovering after hardware failure

Objective	Tasks
“Refreshing the snapshot set on the current disks” on page 153	<ul style="list-style-type: none"> ■ Reattaching healthy snapshot volumes ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Creating the refreshed snapshot set
“Moving the production volumes to different disks and refreshing the snapshot set” on page 160	<ul style="list-style-type: none"> ■ Reattaching healthy snapshot volumes ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Adding mirrors to volumes whose drive letters or mount points were reassigned ■ Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned ■ Creating the refreshed snapshot set

Reviewing the prerequisites

Ensure that you have fulfilled the prerequisites cited below before proceeding with the restore operation.

Observe the following prerequisites:

- Make sure to correct the hardware or software issues and repair the failed volumes prior to using the procedures described below. Note that the `vxsnap restore -r` recovery method should be used only after other options have proved unsuccessful.
 See the Troubleshooting section of the *Veritas Storage Foundation Administrator’s Guide* for more information.
- The complete snapshot set including the corresponding XML metadata file must be available before you perform the restore operation.

Reviewing the sample configuration

Table 9-3 shows the objects that relate to an Exchange storage group, SG2, and are used to illustrate the tasks for hardware recovery.

Table 9-3 Objects for the Exchange storage group

Volume	Drive letter or mount point	Object
SG2-tlogs	I:\Logs	transaction logs volume

Table 9-3 Objects for the Exchange storage group

Volume	Drive letter or mount point	Object
SG2-DB1	G:	mailbox store database volume
SG2-DB2	H:	mailbox store database volume

The storage group SG2 has an associated snapshot set with a metadata file TestDB.xml and the volumes shown in [Table 9-4](#).

Table 9-4 Objects for storage group SG2 snapshot set

Volume	Drive letter or mount point	Object
SG2-tlogssnap	not assigned	snapshot volume of SG2-tlogs
SG2-DB1snap	not assigned	snapshot volume of SG2-DB1
SG2-DB2snap	not assigned	snapshot volume of SG2-DB2

Scenario I: Database and transaction logs volumes are missing

If the metadata file is available but the transaction logs volume and database volume are missing, you can use the snapshot set to restore the storage group to the point in time that the snapshot set was created or last refreshed.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. See [“Identifying the missing volumes”](#) on page 131.
- Use Exchange to dismount all remaining databases in the storage group. See [“Dismounting the Exchange database”](#) on page 131.
- Delete the missing volumes from Storage Foundation for Windows. See [“Deleting missing volumes from Storage Foundation for Windows”](#) on page 133.
- Replace the failed hardware and add the new disks to the dynamic disk group.

Scenario I: Database and transaction logs volumes are missing

See “[Replacing hardware and adding disks to the dynamic disk group](#)” on page 133.

- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.

See “[Changing the drive letter or mount points of the snapshot volumes](#)” on page 134.

- Use the `vxsnap restore` to recover the databases.

See “[Restoring the storage group to the point in time](#)” on page 136.

- Refresh the snapshot set.

See “[Refreshing the snapshot set](#)” on page 137.

Identifying the missing volumes

Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. The drive letter or mount point information is available either through the VEA console or Exchange.

In Exchange, the transaction logs path is shown on the properties page of the storage group and the database path is shown on the properties page of each individual database (store).

For example, if SG2-tlogs and SG2-DB1 are missing due to hardware failure, you would need the information in [Table 9-5](#).

Table 9-5 Information for missing volumes

Status	Volume	Drive letter or mount point	Object
Missing	SG2-tlogs	I:\Logs	transaction logs volume
Missing	SG2-DB1	G:	mailbox store database volume
Available	SG2-DB2	H:	mailbox store database volume
Available	SG2-tlogssnap	not assigned	snapshot volume of SG2-tlogs
Available	SG2-DB1snap	not assigned	snapshot volume of SG2-DB1
Available	SG2-DB2snap	not assigned	snapshot volume of SG2-DB2

Dismounting the Exchange database

Dismount all the databases in the Exchange storage group.

To dismount an Exchange 2003 database

- 1 In Exchange System Manager, navigate to the database that you want to dismount.
- 2 Right-click on the database and click **Dismount Store**.
- 3 Repeat for all databases in the Exchange storage group.

To dismount an Exchange 2007 database

- 1 In the Exchange Management Console, navigate to the database that you want to dismount.
- 2 Right-click the database, and click **Dismount Database**.
- 3 A warning appears asking if you want to dismount the database. Click **Yes**.
- 4 Verify that the status indicated in the Status column has changed from Dismounted to Mounted.
- 5 Repeat the procedure for all databases in the Exchange storage group.

Deleting missing volumes from Storage Foundation for Windows

You must delete missing volumes from Storage Foundation for Windows. This makes the drive letter or mount point available and removes information about the missing volume that is no longer relevant.

Note: Before deleting the missing production volumes, note the volume name and drive letter or mount point. You will need this information later on in order to assign the same drive letter to the snapshot volume associated with each missing production volume.

To delete a missing volume using the VEA console

- 1 Right-click on the missing volume and select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete the volume.

To delete a missing volume using the command line

- ◆ Type the command as in the following example, which deletes the SG2-DB1 volume:

```
> vxassist -gExch1SG2 delete SG2-DB1
```

For complete syntax of `vxassist delete` command, refer to the *Veritas Storage Foundation Administrator's Guide*.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk.

Note the new disk number(s). You will need the information later on to add the disks to the dynamic disk group and for the recovery operation.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 In the Actions menu, click **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).

or

If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console and is marked with (No Signature), for example, `harddisk5 (No signature)`. Proceed to the next step.

Scenario I: Database and transaction logs volumes are missing

- 4 Right-click on a new, unsigned disk and click **Write Signature**.
- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange storage group. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 In the Welcome panel, click **Next**.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Changing the drive letter or mount points of the snapshot volumes

If the production volume is missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

[Table 9-6](#) shows the changes that would be made in the sample configuration.

Table 9-6 Changes to make in the sample configuration

Volume	Drive letter or mount point	Object
SG2-tlogssnap	I:\Logs	snapshot volume of SG2-tlogs
SG2-DB1snap	G:	snapshot volume of SG2-DB1
SG2-DB2snap	Do not change, production volume is healthy	snapshot volume of SG2-DB2

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To change a snapshot volume drive letter to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.

To change a snapshot volume mount point to a production volume drive letter

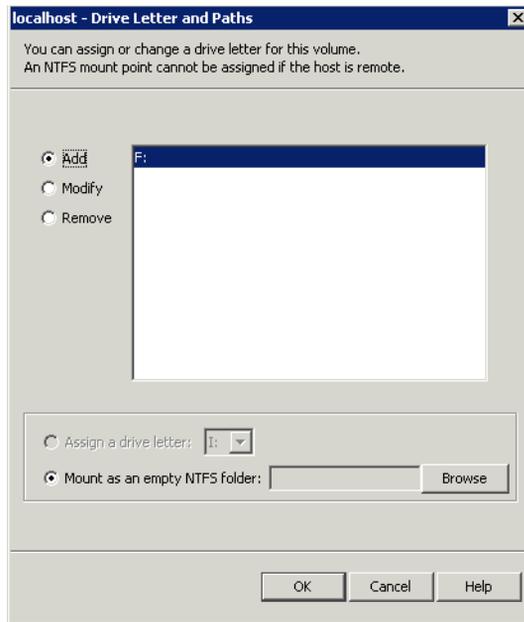
- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Click **Mount as an empty NTFS folder** and click **Browse**.

Scenario I: Database and transaction logs volumes are missing

- 8 Double-click the volume where the production volume was mounted.



- 9 Select the drive folder where the production volume was mounted.
- 10 Click **OK** to assign the mount point.

Restoring the storage group to the point in time

Restore the storage group to the point in time and mount the Exchange databases.

When restoring after hardware failure, you must use the command line and make sure to specify the `-r` option.

See “[vxsnap restore](#)” on page 176.

Caution: You must dismount all databases (stores) in the affected storage group before issuing the command. You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).

To restore the storage group to the point in time of the snapshot set

- 1 Verify that all databases in the storage group are dismounted.

- 2 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).
- 3 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 4 Click **Start > Run**. In the Command window, type the following:

```
vxsnap -x <Filename> [-f] [-b] -r [-a] restorerestoreType=<PIT>
writer=WriterName
```

For example:

```
vxsnap -x TestDB.xml -r -a restore restoreType=PIT writer=
"Microsoft Exchange Writer"
```

This command will restore the storage group associated with the TestDB.xml metadata file to the point in time that the snapshot set was created or last refreshed.

Use the `-a` attribute to dismount and mount the database automatically or use Exchange to dismount and mount the database before and after the restore operation respectively.

For a VCS cluster setup, specifying the `-a` option offlines the VCS resource for the database before a restore operation and onlines the resource after a restore operation.
- 5 Use Exchange to mount the restored databases in the Exchange storage group.

Refreshing the snapshot set

After recovering from a hardware failure, refresh the snapshot set.

See [“Refreshing the snapshot set”](#) on page 153.

In the example, the SG2-tlogssnap and SG2-DB1snap volumes will reside on the disks where the original snapshot volumes were and the SG2-DB2 volume will reside on the disk(s) where the original production volume was.

Scenario II: Database volumes missing, transaction logs are available

If all the database volumes are missing but the transaction logs volume and metadata file are available, you can either use the snapshot set to restore the storage group to the Point in Time (PIT) that the snapshot set was created or last refreshed or you can use the snapshot set to perform a roll-forward recovery to the Point of Failure (POF).

The tasks for either a point in time or point of failure recovery are the same except for the actual `vxsnap restore` command.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
See [“Identifying the missing volumes”](#) on page 138.
- Use Exchange to dismount all remaining databases in the storage group.
See [“Dismounting Exchange databases”](#) on page 139.
- Delete the missing volumes from Storage Foundation for Windows.
See [“Deleting missing volumes from Storage Foundation for Windows”](#) on page 139.
- Replace the failed hardware and add the new disks to the dynamic disk group.
See [“Replacing hardware and adding disks to the dynamic disk group”](#) on page 140.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
See [“Changing the drive letter or mount points of the snapshot volumes”](#) on page 141.
- Use the `vxsnap restore` to recover the databases.
See [“Restoring the storage group to the Point in Time \(PIT\)”](#) on page 143.
See [“Recovering the storage group to the Point of Failure \(POF\)”](#) on page 144.
- Refresh the snapshot set.
See [“Refreshing the snapshot set”](#) on page 152.

Identifying the missing volumes

Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. The drive letter or mount point information is available either through the VEA console or Exchange.

For example, if SG2-DB1 and SG2-DB2 are missing due to hardware failure, you would need the information in [Table 9-7](#).

Table 9-7 Information to identify missing volumes

Status	Volume	Drive letter or mount point	Object
Available	SG2-tlogs	I:\Logs	transaction logs volume

Table 9-7 Information to identify missing volumes

Status	Volume	Drive letter or mount point	Object
Missing	SG2-DB1	G:	mailbox store database volume
Missing	SG2-DB2	H:	mailbox store database volume
Available	SG2-tlogssnap	not assigned	snapshot volume of SG2-tlogs
Available	SG2-DB1snap	not assigned	snapshot volume of SG2-DB1
Available	SG2-DB2snap	not assigned	snapshot volume of SG2-DB2

Dismounting Exchange databases

Dismount all the databases in the Exchange storage group.

To dismount an Exchange 2003 database

- 1 In Exchange System Manager, navigate to the database that you want to dismount.
- 2 Right-click on the database and click **Dismount Store**.
- 3 Repeat for all databases in the Exchange storage group.

To dismount an Exchange 2007 database

- 1 In the Exchange Management Console, navigate to the database that you want to dismount.
- 2 Right-click the database, and click **Dismount Database**.
- 3 A warning appears asking if you want to dismount the database. Click **Yes**.
- 4 Verify that the status indicated in the Status column has changed from Dismounted to Mounted.
- 5 Repeat the procedure for all databases in the Exchange storage group.

Deleting missing volumes from Storage Foundation for Windows

You must delete missing volumes from Storage Foundation for Windows. This will make the drive letter or mount point available and will remove information

about the missing volume that is no longer relevant. For example delete the SG2-DB1 and SG2-DB2 volumes. Do not delete the SG2-tlogs volume.

Note: Before deleting the missing production volumes, note the volume name and drive letter or mount point. You will need this information later on in order to assign the same drive letter to the snapshot volume associated with each missing production volume.

To delete a missing volume using the VEA console

- 1 Right-click on the designated volume and select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete the volume.

To delete a missing volume using the command line

Type the command as in the following example, which deletes the SG2-DB1 volume:

```
vxassist -gExch1SG2 delete SG2-DB1
```

For the complete syntax of the `vxassist delete` command, see *Veritas Storage Foundation Administrator's Guide*.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk. Note the new disk number(s). You will need the information to add the disks to the dynamic disk group and for the recovery operation.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 From the **Actions** menu, select **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).

or

If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console and is marked with (No Signature), for example, `harddisk5 (No signature)`. Proceed to the next step.

- 4 Right-click on a new, unsigned disk and click **Write Signature**.

- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange storage group. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 Click **Next** at the welcome screen.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Changing the drive letter or mount points of the snapshot volumes

If the production volume is missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

[Table 9-8](#) shows the changes that would be made in the sample configuration.

Table 9-8 Changing the drive letter or mount points

Volume	Drive letter or mount point	Object
SG2-tlogsnap	Do not change, production volume is healthy	snapshot volume of SG2-tlogs
SG2-DB1snap	G:	snapshot volume of SG2-DB1
SG2-DB2snap	H:	snapshot volume of SG2-DB2

Note: Assign drive letters to database snapshot volumes only. Do not change the transaction logs volume.

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To change a snapshot volume drive letter to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.
- 5 Repeat this procedure for all snapshot volumes of database volumes only.

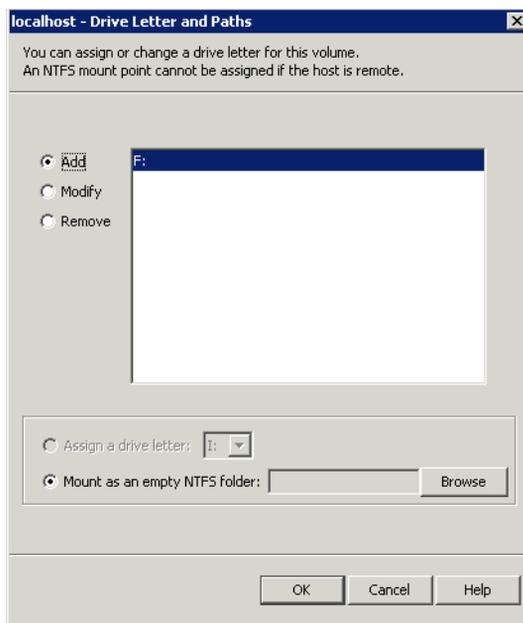
To change a snapshot volume mount point to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.
- 9 Repeat this procedure for all snapshot volumes of database volumes only.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.

- 7 Click **Mount as an empty NTFS folder** and click **Browse**.
- 8 Double-click the volume where the production volume was mounted.



- 9 Select the drive folder where the production volume was mounted.
- 10 Click **OK** to assign the mount point.
 Repeat this procedure for all snapshot volumes of database volumes only.

Restoring the storage group to the Point in Time (PIT)

Restore the storage group to the point in time that the snapshot set was created or last refreshed and mount the Exchange databases.

Refer to Microsoft Exchange Shell cmdlets for details.

Note: When restoring after hardware failure, you must use the command line and make sure to specify the `-r` option.

Caution: If you choose to restore the storage group to the point in time, you cannot later restore it to the point of failure. You can only perform one recovery procedure on a database.

Caution: You must dismount all databases (stores) in the affected storage group before issuing the command. You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).

To restore the storage group to the point in time of the snapshot set

Verify that all databases in the storage group are dismounted.

- 1 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).
- 2 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 3 Click **Start > Run**. In the Command window, type the following: Type the command as in the following example:

```
vxsnap -x <Filename> [-f] [-b] -r [-a] restoreRestoreType=<PIT>  
writer=WriterName
```

For example:

```
vxsnap -x TestDB.xml -r -a restore restoreType=PIT  
writer="Microsoft Exchange Writer"
```

This command will restore the storage group associated with the TestDB.xml metadata file to the point in time that the snapshot set was created or last refreshed.

- 4 Use Exchange to mount the restored databases (stores) in the Exchange storage group.

Recovering the storage group to the Point of Failure (POF)

Use the following procedure to recover the storage group to the point of failure. Refer to the Microsoft Exchange Shell cmdlets.

Note: When restoring after a hardware failure, you must use the command line and make sure to specify the `-r` option.

Caution: If you choose to restore the storage group to the point of failure, you cannot later restore to the point of time. You can only perform one recovery procedure on a storage group.

Caution: You must dismount all databases (stores) in the affected storage group before issuing the command. You must verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).

To recover the storage group to the point of failure

Verify that all databases (stores) in the affected Exchange storage group are dismounted.

- 1 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).
- 2 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 3 Click **Start > Run**. In the Command window, type the following command as which recovers the storage group associated with the TestDB.xml metadata file:

```
vxsnap -x TestDB.xml -r restore restoreType=POF  
writer="Microsoft Exchange Writer"
```
- 4 Mount all databases (stores) in the Exchange storage group; the transaction logs will roll forward to the point of failure. This action can be observed in the Event Viewer Application log.

Refreshing the snapshot set

After recovering from a hardware failure, refresh the snapshot set.

See “[Refreshing the snapshot set](#)” on page 153.

In the example, SG2-DB1snap and SG2-DB2snap will reside on the disks where the original snapshot volumes were and SG2-tlogs will reside on the disk(s) where the original production volume was.

Scenario III: Some DB volumes missing, transaction logs are available

If some but not all of the database volumes are missing and the transaction logs volume and metadata file are available, you can use the snapshot set to perform a roll-forward recovery to the point of failure of the failed volumes. If only one database volume is missing, you can perform a single database restore by using the `subComponent=subComponentName` attribute in the `vxsnap restore` command.

Complete the following tasks to perform a VSS-integrated recovery:

- Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
See [“Identifying the missing volumes”](#) on page 146.
- Use Exchange to dismount all remaining databases in the mailbox database storage group.
See [“Dismounting Exchange databases”](#) on page 147.
- Delete the missing volumes from Storage Foundation for Windows.
See [“Deleting missing volumes from Storage Foundation for Windows”](#) on page 148.
- Replace the failed hardware and add the new disks to the dynamic disk group.
See [“Replacing hardware and adding disks to the dynamic disk group”](#) on page 148.
- Reassign the drive letters or mount points of the snapshot volumes so that they are the same as the missing production volumes.
See [“Changing the drive letter or mount points of the snapshot volumes”](#) on page 149.
- Use the `vxsnap restore` to recover the databases.
See [“Recovering the storage group to the Point of Failure \(POF\)”](#) on page 151.
- Refresh the snapshot set.
See [“Refreshing the snapshot set”](#) on page 152.

Identifying the missing volumes

Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume. The drive letter or mount point information is available either through the VEA console or Exchange. In Exchange, the transaction logs path is shown on the properties page of the storage group, and the database path is shown on the properties page of each individual database (store).

For example, if SG2-DB1 is missing due to hardware failure, you would need the information in [Table 9-9](#).

Table 9-9 Information for missing SG2-DB1

Status	Volume	Drive letter or mount point	Object
Available	SG2-tlogs	I:\Logs	transaction logs volume
Missing	SG2-DB1	G:	mailbox store database volume
Available	SG2-DB2	H:	mailbox store database volume
Available	SG2-tlogssnap	not assigned	snapshot volume of SG2-tlogs
Available	SG2-DB1snap	not assigned	snapshot volume of SG2-DB1
Available	SG2-DB2snap	not assigned	snapshot volume of SG2-DB2

Dismounting Exchange databases

Dismount all the databases in the Exchange storage group.

To dismount an Exchange 2003 database

- 1 In Exchange System Manager, navigate to the database that you want to dismount.
- 2 Right-click on the database and click **Dismount Store**.
- 3 Repeat for all databases in the Exchange storage group.

To dismount an Exchange 2007 database

- 1 In the Exchange Management Console, navigate to the database that you want to dismount.
- 2 Right-click the database, and click **Dismount Database**.
- 3 A warning appears asking if you want to dismount the database. Click **Yes**.
- 4 Verify that the status indicated in the Status column has changed from Dismounted to Mounted.
- 5 Repeat the procedure for all databases in the Exchange storage group.

Deleting missing volumes from Storage Foundation for Windows

You must delete missing volumes from Storage Foundation for Windows. This will make the drive letter or mount point available and will remove information about the missing volume that is no longer relevant.

For example delete the SG2-DB1 volume. Do not delete the SG2-DB2 or SG2-tlogs volumes.

Note: Before deleting the missing production volumes, note the volume name and drive letter or mount point. You will need this information later on for assigning the same drive letter to the snapshot volume associated with each missing production volume.

To delete a missing volume using the VEA console

- 1 Right-click on the designated volume and select **Delete Volume**.
- 2 You are prompted for verification. Click **Yes** to delete the volume.

To delete a missing volume using the command line

Type the command as in the following example, which deletes the SG2-DB1DB1 volume:

```
vxassist -gExch1SG2 delete SG2-DB1
```

For the complete syntax of the `vxassist delete` command, refer to the *Veritas Storage Foundation Administrator's Guide*.

Replacing hardware and adding disks to the dynamic disk group

Replace any defective hardware and add new disks to the dynamic disk group, as necessary. The number assigned to a new disk, for example `harddisk5`, may not be the same as the disk number of the failed disk. Note the new disk number(s). You will need the information to add the disks to the dynamic disk group and for the recovery operation.

To replace the hardware and add the new disks to the dynamic disk group

- 1 Replace the defective hardware.
- 2 From the **Actions** menu, select **Rescan**.
- 3 If the disk was previously used in another system and has a disk signature, proceed to [step 7](#).

or

If the new disk has never been used before, it is unsigned and needs a disk signature. In this case, the disk appears in the left pane of the VEA console

and is marked with (No Signature), for example, harddisk5 (No signature). Proceed to the next step.

- 4 Right-click on a new, unsigned disk and click **Write Signature**.
- 5 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 6 Click **OK**.
After a signature appears on a disk, the disk will display as a basic disk.
- 7 Add the disk to the dynamic disk group of the volumes associated with the production Exchange storage group. Right-click the new disk and click **Add Disk to Dynamic Disk Group**.
- 8 Click **Next** at the welcome screen.
- 9 Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
- 10 Click **Next**.
- 11 Review the confirmation information and click **Next**.
- 12 Click **Finish** to upgrade the selected disks from basic to dynamic and add them to the dynamic disk group.

Changing the drive letter or mount points of the snapshot volumes

If the production volume is missing, change the drive letter or mount point of the snapshot volume to the drive letter or mount point that was assigned to the missing production volume. If the production volume is healthy and available, do not make any changes.

[Table 9-10](#) shows the changes that would be made in the sample configuration.

Table 9-10 Changes if the production volume is missing

Volume	Drive letter or mount point	Object
SG2-tlogssnap	Do not change, production volume is healthy	snapshot volume of SG2-tlogs
SG2-DB1snap	G:	snapshot volume of SG2-DB1
SG2-DB2snap	Do not change, production volume is healthy	snapshot volume of SG2-DB2

Note: Assign drive letters or mount points to database snapshot volumes of missing production volumes only. Do not change the drive letter or mount point for healthy database volumes or for the transaction logs volume.

The steps for changing a drive letter vary slightly from the steps for changing a mount point. Follow the procedure that best fits your environment.

To change a snapshot volume drive letter to a production volume drive letter

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Select **Modify**.
- 3 From the assign drive letter list, select the drive letter originally assigned to the production volume.
- 4 Click **OK**.
- 5 Repeat this procedure for all snapshot volumes of database volumes only.

To change a snapshot volume mount point to a production volume drive letter

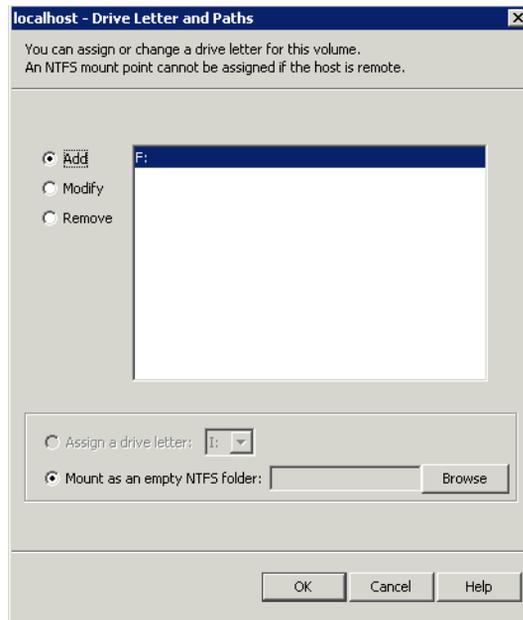
- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.
- 4 Click **Yes** to confirm your choice.
- 5 Assign the new drive letter. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Select the drive letter originally assigned to the associated production volume.
- 8 Click **OK**.
- 9 Repeat this procedure for all snapshot volumes of database volumes only.

To change a snapshot volume mount point to a production volume mount point

- 1 Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 2 Click **Remove**.
- 3 Click **OK**.

Scenario III: Some DB volumes missing, transaction logs are available

- 4 Click **Yes** to confirm your choice.
- 5 Assign the new mount point. Right-click the snapshot volume, click **File System** and click **Change Drive Letter and Path**.
- 6 Click **Add**.
- 7 Click **Mount as an empty NTFS folder** and click **Browse**.
- 8 Double-click the volume where the production volume was mounted.



- 9 Select the drive folder where the production volume was mounted.
- 10 Click **OK** to assign the mount point.
Repeat this procedure for all snapshot volumes of missing database volumes only.

Recovering the storage group to the Point of Failure (POF)

Use the following procedure to recover the mailbox database storage group to the point of failure.

Note: When restoring after hardware failure, you must use the command line and make sure to specify the `-r` option.

Caution: You must dismount all databases (stores) in the affected mailbox database storage group before issuing the command. You must also verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the mailbox database storage group (component).

To recover the storage group to the point of failure from the command line

- 1 Verify that all databases (stores) in the affected Exchange mailbox database storage group are dismounted.
- 2 Verify that you have correctly assigned the drive letter or mount point to each volume and that you have accounted for all the volumes in the storage group (component).
- 3 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 4 Type the command as in the following example, which recovers the mailbox database storage group associated with the TestDB.xml metadata file:

```
vxsnap -x TestDB.xml -r restore restoreType=POF  
writer="Microsoft Exchange Writer"
```

If only one database volume is missing, you can perform a single database restore by using the `subComponent=subComponentName` option in the `vxsnap restore` command. For example, if DB1 is the name of the Exchange database that is associated with the volume SG2-DB1, type the command:

```
> vxsnap -x TestDB.xml -r restore restoreType=POF  
subComponent=DB1 writer="Microsoft Exchange Writer"
```

- 5 Use Exchange to mount all databases (stores) in the Exchange mailbox database storage group; the transaction logs will roll forward to the point of failure.

This action can be observed in the Event Viewer Application log.

Refreshing the snapshot set

After recovering from a hardware failure, refresh the snapshot set.

See [“Refreshing the snapshot set”](#) on page 153.

In the example, SG2-DB1snap will reside on the disk(s) where the original snapshot volume was and the SG2-DB2 and SG2-tlogs volumes will reside on the disk(s) where the original production volume was.

Refreshing the snapshot set

Refresh your snapshot set after recovering from the hardware failure. You must decide if the disks which now store the data are suitable for production or if you want to move the production volumes to different disks which you earlier added to the dynamic disk group as replacement disks.

Note: Note that during the process of refreshing the snapshot set any volumes that were originally snapshot volumes of missing production volumes will become master (production) volumes. These volumes currently reside on the disks on which the original snapshot volumes were created.

Choose one of the following options for your snapshot set:

- [“Refreshing the snapshot set on the current disks”](#) on page 153.
- [“Moving the production volumes to different disks and refreshing the snapshot set”](#) on page 160.

Refreshing the snapshot set on the current disks

Perform the following tasks to refresh the snapshot set on the current disks. In this case the volumes that were originally snapshot volumes of missing production volumes will become production volumes and will reside on the disks on which the original snapshot volumes were created.

Complete the following tasks to refresh the snapshot set on the current disks:

- Snap back the snapshot transaction logs volumes if you recovered using a roll-forward recovery to the point of failure. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them), as might occur in Scenario III.
See [“Reattaching healthy snapshot volumes”](#) on page 154.
- Snap clear each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. The snap clear removes the association of the snapshot volume with the original production volume.
See [“Clearing the snapshot association for volumes whose drive letters or mount points were reassigned”](#) on page 155.
- Prepare the snapshot mirror for each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process.
See [“Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned”](#) on page 155.

- Use the `vxsnap create` command to create a new snapshot set of all the volumes in the storage group.
See “[Creating the new snapshot set](#)” on page 156.

Reattaching healthy snapshot volumes

If you recovered using a roll-forward recovery to the Point of Failure (POF), you must snap back the snapshot transaction logs volume using the **Resynchronize using the original volume** option. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them), as might occur in Scenario III.

The VSS Snapback wizard reattaches and resynchronizes an existing snapshot snapshot set so that it matches the current state of its original Exchange storage group.

For the complete syntax of the `vxassist snapback` command, see *Veritas Storage Foundation Administrator's Guide*.

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 be accessing the snapshot set.
- 2 From the VEA console URL bar, select the `<host name>` which is 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 on the node of the 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 required information needed 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**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close all open handles on the snapshot volumes.
- 2 Type the following command, replacing the example values:

```
> vxassist -gExch1-SG2 snapback SG2-tlogsnap
```

If there is an open handle on the volume you will get the following error message “One of the volumes of this dynamic disk group is currently in use. Please close applications and try again.”
- 3 Repeat the command for any other snapshot volumes in the storage group that were not affected by the disk failure.

Clearing the snapshot association for volumes whose drive letters or mount points were reassigned

The clear removes the association of the snapshot volume with the original production volume. Perform the snap clear operation only on those volumes whose drive letter or mount point was reassigned prior to the restore process. Volumes that remained healthy throughout the disk failure have an existing snapshot mirror that will continue to be used as part of the snapshot set.

To snap clear a volume

- 1 Right-click the snapshot volume and select **Snap > Snap Clear**.
- 2 Click **Yes** to confirm the snap clear operation for the specified volume.
- 3 Repeat this process for each volume whose drive letter or mount point was reassigned prior to the restore process.

Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned

Use the Prepare command to prepare each of the volumes whose drive letter or mount point was reassigned prior to the restore process. Make sure to select disks or LUNs that are not used for production data. However, you may create more than one snapshot volume on the same disk or LUN as long as there is sufficient space available.

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare**.
- 2 In the **Prepare** dialog box:
 - Choose **Manually select disks**.

- Use the **Add** and **Remove** buttons to move the desired disks to the **Selected disks** box.
 - Click **OK**.
- 3 Repeat the procedure to create a snapshot mirror for each volume whose drive letter or mount point was reassigned prior to the restore process. Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before continuing with the `vxsnap create` command.

To create snapshot mirrors from the command line

- 1 Type the command as in the following example, which prepares volume G: and selects `harddisk3` for the split-mirror snapshot volume:

```
vxassist prepare G: Harddisk3
```
- 2 Repeat [step 1](#) to create snapshot mirrors of each volume associated whose drive letter or mount point was reassigned prior to the restore process. Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before continuing with the `vxsnap create` command.

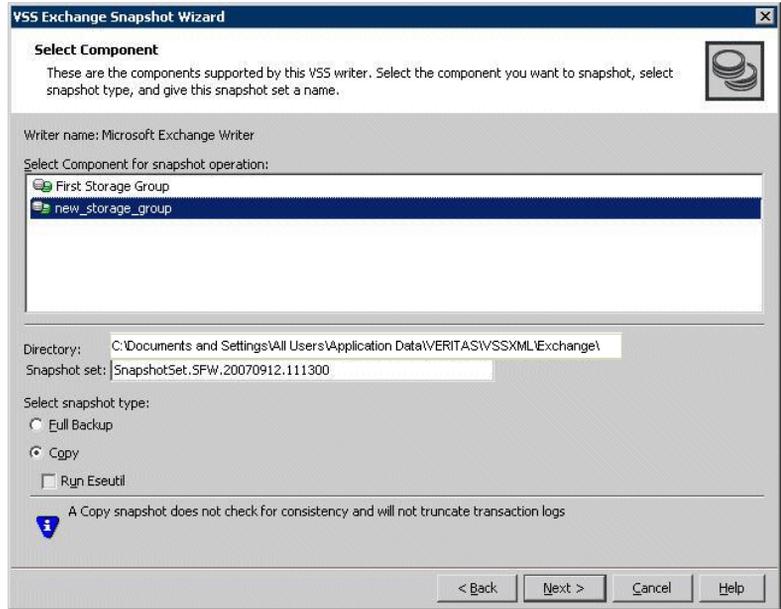
Creating the new snapshot set

Create the new snapshot set from the VEA console or from the command line.

To create the snapshot set from the VEA console

- 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 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange 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**.
- 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 2003 or 2007, the component is a storage
group.

Directory	<p>Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory 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 will fail.</p> <p>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 "redirect.txt." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, <code>G:\BackupSets</code>. Save the redirect.txt file in the default directory <code>C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML</code> on each node of the cluster.</p> <p>Note: 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 will fail.</p>
Snapshot set	<p>Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.</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	<p>In the Change Attributes panel, optionally change the attributes for the snapshot volumes as follows and click Next:</p>
Snapshot Volume Label	<p>Displays the read-only label for the snapshot volume.</p>

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 Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- 8 Review the specifications of the snapshot set and click **Finish**.

To create the new snapshot set from the command line

- ◆ Type the `vxsnap create` command. The exact command will vary depending on your environment.

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<newVolName>] [/Plex=<plexName>] ...writer=<writerName>
component=<componentName> [backuptype=<backuptype>] [-E] [-O] [-C] [
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Any text string that contains spaces must be enclosed in quotation marks.

Note: If you have scripted the refresh process, you will need to change the snapshot volume names assigned by your script using the option `Newvol=NewVolName`. For instance, if SG2-DB1 volume was missing, and the volume SG2-DB1snap was used in the restore process, it will retain the volume name SG2-DB1snap. Your script will not be able to assign SG2-DB1snap as the new snapshot volume name in step 4 when using the `vxsnap create` command.

See “[vxsnap create](#)” on page 172.

Moving the production volumes to different disks and refreshing the snapshot set

During the process of refreshing the snapshot set any volumes that were originally snapshot volumes of missing production volumes will become master (production) volumes. These volumes currently reside on the disks on which the original snapshot volumes were created.

Perform the following tasks to move the master (production) volumes to different disks:

- Snap back the snapshot transaction logs volumes if you recovered using a roll-forward recovery to the point of failure. Additionally, snap back the snapshot volumes of any databases that were not affected by the disk failure (and thus still have snapshot volumes associated with them) as might occur in Scenario III.
See [“Reattaching healthy snapshot volumes”](#) on page 160.
- Snap clear each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. The snap clear removes the association of the snapshot volume with the original production volume.
See [“Clearing the snapshot association for volumes whose drive letters or mount points were reassigned”](#) on page 161.
- Add a mirror to each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. For each new mirror (plex), choose the disks where you want the production volumes to reside.
See [“Adding mirrors to volumes whose drive letters or mount points were reassigned”](#) on page 162.
- Prepare each of the new mirror volumes and convert an existing mirror into a snap plex. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex).
See [“Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned”](#) on page 162.
- Use the `vxsnap create` command to create a new snapshot set of all the volumes in the storage group.
See [“Creating the new snapshot set”](#) on page 163.

Reattaching healthy snapshot volumes

If you recovered using a roll-forward recovery to the point of failure, use the VSS Snapback wizard or `vxsnap reattach` command to reattach the snapshot transaction logs volume and the snapshot volumes of any databases that were

Moving the production volumes to different disks and refreshing the snapshot set

not affected by the disk failure (and thus still have snapshot volumes associated with them), as might occur in Scenario III.

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 be accessing the snapshot set.
- 2 From the VEA console URL bar, select the *<host name>* which is 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 on the node of the 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 required information needed 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**.

To reattach the split-mirror snapshots to the original volumes from the command line

- 1 Close the Exchange GUI and all Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.
- 2 Type the command, replacing the example file name:

```
vxsnap -x Image1.xml reattach writer="Microsoft Exchange Writer"
```

The complete syntax for the `vxsnap reattach` command is:

```
vxsnap -x <filename> [-f][-b] reattach writer=<writername>  
[secHosts=<secondary hosts>]
```

Clearing the snapshot association for volumes whose drive letters or mount points were reassigned

The `snap clear` removes the association of the snapshot volume with the original production volume. Perform the `snap clear` operation only on those volumes

whose drive letter or mount point was reassigned prior to the restore process. Volumes that remained healthy throughout the disk failure have an existing snapshot mirror that will continue to be used as part of the snapshot set.

To snap clear a volume

- 1 Right-click the snapshot volume and select **Snap > Snap Clear**.
- 2 Click **Yes** to confirm the Snap Clear operation for the specified volume.
- 3 Repeat this process for each volume whose drive letter or mount point was reassigned prior to the restore process.

Adding mirrors to volumes whose drive letters or mount points were reassigned

Add a mirror to each volume in the snapshot set whose drive letter or mount point was reassigned prior to the restore process. For each new mirror (plex), choose the disks where you want the production volumes to reside.

To add a mirror to a volume

- 1 Right-click on the volume you want to mirror, click **Mirror** and click **Add**.
- 2 Enter the number of mirrors to create.
- 3 Choose to manually assign the destination disks for the mirrors.
- 4 Select the disks in the Available disks list where you want the production volume to reside, and use the **Add** button to move them to the Selected disks list.
- 5 Click **OK** to add the mirror.

Note: 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.

Creating snapshot mirrors of volumes whose drive letters or mount points were reassigned

Prepare each of the new mirror volumes and choose to convert an existing mirror into a snap plex. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex).

Moving the production volumes to different disks and refreshing the snapshot set

To create the snapshot mirrors using the VEA console

- 1 Right-click the desired volume, select **Snap > Snap Prepare**.
- 2 The result of this command varies depending on whether or not there are additional disks available to create a new mirror. Choose from the following:

If the volume is mirrored and no additional disks are available to create a new mirror

- Click the mirror where you want the snapshot volume to be created. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex). The remaining mirror will become the master (production) volume.
- Click **OK**.

If the volume is mirrored and there are additional disks available on your system

Choose an existing mirror for the snapshot

- Click **Select existing mirror for snap**.
- Click the mirror where you want the snapshot volume to be created. Choose the mirror that was the snapshot volume of the failed production volume as the new snapshot mirror (snap plex). The remaining mirror will become the master (production) volume.
- Click **OK**.

- 3 Repeat the procedure to create a snapshot mirror for each volume associated with the Exchange storage group databases and transaction logs.

Caution: Make sure that the lower pane of the VEA console shows that the resynchronization process is complete before creating the new snapshot set.

Creating the new snapshot set

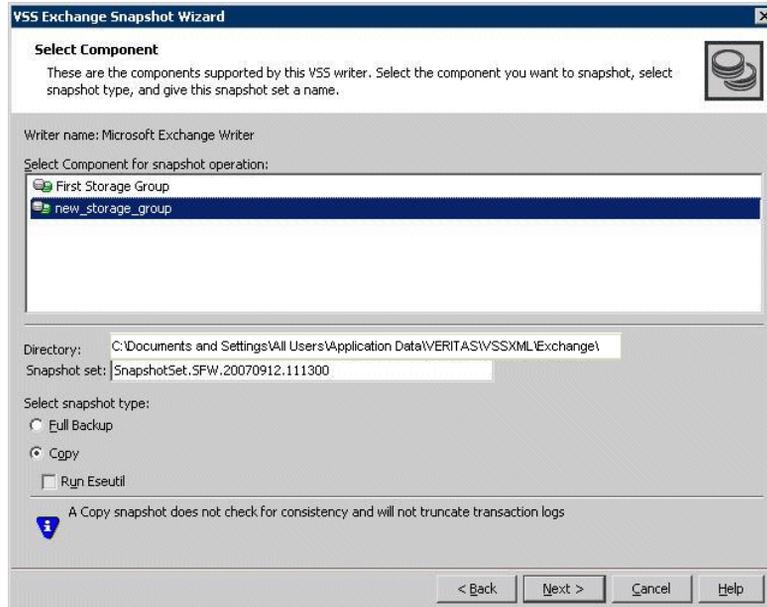
Create the new snapshot set from the VEA console or from the command line.

To create the snapshot set from the VEA console

- 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 Choose one of the following:
 - Right-click **Exchange** and click **VSS Exchange Snapshot**.

Moving the production volumes to different disks and refreshing the snapshot set

- (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**.
- 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 2003 or 2007, the component is a storage group.

Directory

Enter a directory location for the XML file or accept the default. The XML file is stored by default in the directory shown on the screen.

Note: The XML file for the snapshot must be stored separately from the volumes that are included in the snapshots, otherwise a restore will fail.

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 "**redirect.txt**." This text file should contain a single text line specifying the full path to the location of the metadata file, for example, G:\BackupSets. Save the **redirect.txt** file in the default directory C:\Program Files\Veritas\Veritas Volume Manager 5.1\VSSXML on each node of the cluster.

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 will fail.

Snapshot set

Enter a name for the snapshot set. The snapshot set metadata XML file is stored under this name.

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 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.

- 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 Veritas Volume Replicator (VVR). Otherwise, you will be directly taken to the Schedule Information panel.

- Review the specifications of the snapshot set and click **Finish**.

To create the new snapshot set from the command line

- Type the `vxsnap create` command. The exact command will vary depending on your environment. For example:

```
> vxsnap -x TestDB.xml create source=G:/Newvol=SG2-DB1snap2
source=H:/Newvol=SG2-DB2snap2 source=I:/Newvol=SG2-tlogsnap2
writer="Microsoft Exchange Writer" component=SG2
backuptype=copy
```

The complete syntax of the `vxsnap create` command is:

```
vxsnap [-x <filename>] create [source=<volume>]
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>] [/Newvol=<n
ewVolName>] [/Plex=<plexName>] ...writer=<writerName>
component=<componentName> [backuptype=<backuptype>] [-E] [-O] [-C] [
secHosts=<secondary hosts>]
```

The *WriterName* and *ComponentName* must be included in the command. The option to assign drive letters or mount points is useful for tracking volumes and for scripting purposes. Any text string that contains spaces must be enclosed in quotation marks.

Note: If you have scripted the refresh process, you will need to change the snapshot volume names assigned by your script using the option `Newvol=NewVolName`. For instance, if SG2-DB1 volume was missing, and the volume SG2-DB1snap was used in the restore process, it will retain the volume name SG2-DB1snap. Your script will not be able to assign SG2-DB1snap as the new snapshot volume name in step 4 when using the `vxsnap create` command.

See [“vxsnap create”](#) on page 172.

Vxsnap utility command line reference

This chapter includes the following topics:

- [About the vxsnap utility](#)
- [Vxsnap keywords](#)

About the vxsnap utility

The command line utilities are available in the Veritas Storage Foundation for Windows installation directory.

The vxsnap utility integrates with the Windows Volume Shadow Copy Service (VSS) as a VSS Requester. This allows for the simultaneous snapshot of all volumes associated with an Exchange Server database.

Note the following requirements:

- Disk groups must be of a Storage Foundation for Windows 4.0 or later version. You must upgrade any disk groups created using an earlier version of Volume Manager for Windows before using the vxsnap utility
- The CLI commands run only on the server. They will not run on the Veritas Storage Foundation for Windows client.
- The vxsnap commands must be invoked on a local system.
- For Windows Server 2008, all CLI commands must run in the command window in the “run as administrator” mode.

Vxsnap keywords

Type the following sequence to view a description of the command syntax:

vxsnap keyword -?

vxsnap has the following keywords:

prepare	Creates snapshot mirrors of the volumes in the specified component. The snapshot mirrors remain attached to and synchronized with the original volumes Either the <code>prepare</code> or <code>start</code> keyword may be used in the CLI; however <code>prepare</code> is recommended.
create	Creates simultaneous snapshots of all volumes in the specified component, providing a point-in-time snapshot set.
reattach	Reattaches and resynchronizes an existing snapshot set to the original database volumes.
restore	Restores an Exchange storage group (component) or a single database (subcomponent) from a snapshot set. Storage group recovery can be either a point-in-time restore or a roll-forward recovery to the point of failure. Database recovery must be a roll-forward recovery to the point of failure.

vxsnap prepare

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.

Syntax

```
vxsnap prepare component=<componentName>/writer=<writerName>  
[-b] [source=<volume>/harddisk=<harddisk,...>] ...]
```

Attributes

The following attributes apply:

<code>component=<ComponentName></code>	Name of the component; for Exchange 2003 or 2007, this is the storage group name, for example, "SG2".
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer. For Exchange 2007, if local continuous replication (LCR) is enabled and you want to create a snapshot mirror for the passive copy, specify the writer option as "Microsoft Exchange Writer Replica." That is the name of the writer used for the passive copy.
<code>-b</code>	Resynchronizes the volume in the background. A snapshot cannot be made until the resynchronization is complete.
<code>source=<Volume></code>	Indicates the source volume for the snapshot mirror 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 prepared, for example, <code>harddisk2</code> .

Example

```
vxsnap prepare component=SG2/writer="Microsoft Exchange Writer"  
source=L:/harddisk=harddisk2 source=M:/harddisk=harddisk3
```

This command will create snapshot mirrors of all the volumes contained in the Exchange storage group (component) SG2. The snapshot mirror of the volume mounted on L: will be created on disk 2 and the snapshot mirror of the volume mounted on M: will be created on disk 3.

vxsnap create

Creates simultaneous snapshots of all volumes in a specified component, providing a point-in-time snapshot set.

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.

See [“Creating a one-time snapshot set”](#) on page 104.

Syntax

```
vxsnap [-x <filename>] create [source=<volume>]  
[/DriveLetter=<driveLetter>] [/DrivePath=<drivePath>]  
[/Newvol=<newVolName>] [/Plex=<plexName>]...writer=<writerName>  
component=<componentName>[backuptype=<backuptype>] [-E] [-O] [-C]  
[secHosts=<secondary hosts>]
```

Attributes

The following attributes apply:

- | | |
|----------------------------|---|
| <p>-x <Filename></p> | <p>Indicates the name to be assigned to the XML metadata file that will be created with the command. The file name must include the ".xml" extension. The default path to the file is under the SFW program files directory (normally C:\Documents and Settings\All Users\Application Data\Veritas\VSSXML\Exchange). If you wish to place the file in another directory, specify a full path before the file name, for example J:\XML\Image1.xml.</p> |
|----------------------------|---|

<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.
<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>	Specifies the name of the new snapshot volume that is created. If the name is not specified using this option, the form "SnapVolume01" is created. The full device path becomes: \Device\HarddiskDmVolumes\ <DiskGroupName>\<NewVolName>
<code>writer=<WriterName></code>	Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer. If you plan to be able to do a roll-forward recovery to the point of failure, you must specify the writer. For Exchange 2007, if local continuous replication (LCR) is enabled and you want to create a snapshot of the passive copy, specify the writer option as "Microsoft Exchange Writer Replica."

<code>component=<ComponentName></code>	Name of the component; for Exchange 2003 and 2007, this is the storage group name, for example, SG2. If you plan to be able to do a roll-forward recovery to the point of failure, you must use this component.
<code>backuptype=<Backuptype></code>	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>sechosts=<SecondaryHosts></code>	Applies to Veritas Volume Replicator (VVR) 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 backupdoc.xml create  
writer="Microsoft Exchange Writer"  
component="SG2"  
source=L:/DriveLetter=O source=M:/DriveLetter=P  
source=N:/DriveLetter=Q backuptype=full
```

This example creates a snapshot set based on the component “SG2,” which contains volume L, the log volume, and volumes M and N, two database volumes. The snapshots are assigned drive letters O, P, and Q, respectively. The XML file backupdoc.xml is used to store the VSS metadata that identifies the snapshot set. The Full backup creates the copy of the database and transaction logs volumes, runs Eseutil to check for consistency, and then truncates the transaction logs.

```
vxsnap -x snapdata.xml  
createsource=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 volume 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 reattach

Reattaches and resynchronizes the snapshot volumes in the snapshot set to the original database volumes.

Note: Make sure that the snapshot volumes are not in use before using this command.

Syntax

```
vxsnap -x <filename> [-f] [-b] reattach writer=<writername>  
[secHosts=<secondary hosts>]
```

Attributes

The following attributes apply:

- | | |
|---------------|---|
| -x <Filename> | The file 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 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, "Microsoft Exchange Writer" or the GUID for the writer. If you plan to be able to do a roll-forward recovery to the point of failure, you must specify the writer.
<code>secHosts=<secondary hosts></code>	Applies to Veritas Volume Replicator (VVR) environment only. Comma separated list of secondary hosts on which a synchronized snapshot is to be taken.

Example

```
vxsnap -x snapdata.xml reattach writer="Microsoft Exchange Writer"
```

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

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 the storage group or a single database.

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 Exchange to dismount all the databases in the storage group and then mount them after the command is completed. For Exchange 2007, you have the following additional options:

- If you want to restore from a passive copy snapshot, additional steps are required. See [“Recovery using an Exchange 2007 passive copy snapshot”](#) on page 113 in [Chapter 8, “Recovering Exchange mailbox storage group or databases”](#).
- If you have created a recovery storage group (RSG) for a storage group, you have the option to restore snapshot volumes to the databases in the RSG.

Note: After completing a point of failure (POF) recovery of a single database, Symantec recommends using the **vxsnap reattach** command to reattach and resynchronize the other databases in the storage group and to use the **vxsnap create** command to create a new snapshot set.

Syntax

```
vxsnap -x <filename>[-f] [-b] [-r] [-a] restore  
restoreType=<PIT|POF>writer=<writername> [subComponent=<subComponentName>] [RSG=<Yes|No>]
```

Attributes

The following attributes apply:

- | | |
|-----------------------|--|
| -x <Filename> | The file created by the vxsnap create command. Each snapshot set must have a unique name for the metadata file. |
| -f | Forces the snapback. Make sure 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. |
| -r | Recover one or more of the original volumes that are missing. Example below shows additional required steps. |
| -a | Use the [-a] attribute to dismount and mount the database automatically. |
| restoreType=<PIT POF> | PIT specifies a restore to the point in time that the snapshot set was created or last refreshed. |
| > | POF specifies a roll-forward recovery to the point of failure. |

<code>writer=<WriterName></code>	<p>Unique ID of the VSS writer, for example, "Microsoft Exchange Writer" or the GUID for the writer.</p> <p>If you specify the POF option, you must specify the writer.</p> <p>For Exchange 2007, for which there can be both an active writer and a replica writer, you must specify the active writer "Microsoft Exchange Writer" with the POF option.</p>
<code>subComponent= <subComponentName></code>	<p>Name of the subcomponent to be restored. In Exchange 2003 or 2007, a subcomponent is a mailbox store (database). Use this attribute only in a point of failure recovery.</p>
<code>[RSG=<Yes No>]</code>	<p>Exchange 2007 only.</p> <p>Yes restores the snapshot set to a recovery storage group (RSG). In order to restore to an RSG, the RSG for that storage group must already exist, and databases created inside the RSG must have the same names as that of the source storage group.</p> <p>When performing a PIT recovery for an RSG, all subcomponents of the storage group are recovered.</p> <p>When performing a POF recovery for an RSG, only the subcomponents that you specify in the command are recovered.</p>

Examples

Restore to the Point in Time

```
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 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 subcomponent=DB1  
writer="Microsoft Exchange Writer"
```

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:

Table 10-1 Available recovery type after missing production volume

Scenario	Database Volumes	Transaction Logs Volume	Recovery
Scenario I	One or more volumes are missing.	Missing	point in time
Scenario II	All volumes are missing.	Available	point in time or point of failure
Scenario III	One or more volumes are missing. At least one volume is available,	Available	point of failure

Caution: 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).

It is important that each of the tasks be performed exactly as described.

See [Chapter 9, “Recovering after hardware failure”](#) on page 125 for details.

Complete the following tasks to perform a VSS-integrated recovery:

- Prepare for the recovery
 - Identify the snapshot volume associated with each missing production volume. Note the drive letter or mount point of each volume.
 - Use Exchange Management Console 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.
- Use the `vxsnap restore` command to recover the databases by including the `-r` switch in the `vxsnap restore` command.
For example, run the following command:

```
vxsnap -x snapdata.xml -r 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.
- Refresh the snapshot set.

Microsoft Clustering Solutions

This section includes the following chapters

- [Deploying SFW with MSCS: New Exchange installation](#)
- [Deploying SFW with Microsoft failover clustering: New Exchange installation](#)
- [Deploying SFW with MSCS and Exchange in a campus cluster](#)
- [Deploying SFW with Microsoft failover clustering and Exchange in a campus cluster](#)
- [Deploying SFW and VVR with MSCS: New Exchange installation](#)
- [Deploying SFW and VVR with Microsoft failover clustering: New Exchange installation](#)

About Microsoft clustering solutions

Microsoft clustering may be used with Veritas Storage Foundation for Windows to provide high availability for Exchange.

Microsoft clustering may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide replication support for Exchange. Using VVR with Microsoft clustering provides a replicated backup of your Exchange data, which can be used for recovery after an outage or disaster. However, this

solution does not provide the automated failover capability for disaster recovery that can be achieved using VVR with VCS.

Microsoft clustering solutions are covered in separate chapters according to operating system:

- Microsoft Cluster Server (MSCS) on Windows Server 2003 (Exchange 2003 and 2007)
- Microsoft failover clustering on Windows Server 2008 (Exchange 2007 only)

Deploying SFW with MSCS: New Exchange installation

This chapter includes the following topics:

- [Tasks for a new Exchange installation with SFW and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing an MSCS cluster](#)
- [Creating the MSDTC resource \(Exchange 2003\)](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Preparing the forest and domain \(Exchange 2003\)](#)
- [Adding a Volume Manager Disk Group resource for Exchange 2007 installation](#)
- [Installing Exchange Server](#)
- [Creating an Exchange virtual server group \(Exchange 2003\)](#)
- [Adding Volume Manager Disk Group resources to the Exchange 2007 group](#)
- [Moving Exchange databases and logs to shared storage \(Exchange 2003\)](#)
- [Moving Exchange databases and logs to shared storage \(Exchange 2007\)](#)
- [Implementing a dynamic mirrored quorum resource](#)
- [Verifying the cluster configuration](#)

Tasks for a new Exchange installation with SFW and MSCS (Windows Server 2003)

On Windows Server 2003, you can install and configure Veritas Storage Foundation for Windows with MSCS and Microsoft Exchange Server 2003 or Exchange Server 2007. The environment involves an active/passive configuration with one to one failover capabilities.

Exchange 2007 also runs on Windows Server 2008. For information on deploying SFW with Exchange Server 2007 on a Microsoft failover cluster on Windows Server 2008, see:

- [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#) on page 231

[Table 11-1](#) outlines the high-level objectives and accompanying tasks that you perform, in the typical order in which you perform them.

Table 11-1 Tasks for installing and configuring SFW with MSCS and Exchange

Objective	Tasks
“Reviewing the requirements” on page 186	<ul style="list-style-type: none">■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 188	<ul style="list-style-type: none">■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster■ Reviewing the benefits of a dynamic mirrored quorum
“Configuring the storage hardware and network” on page 190	<ul style="list-style-type: none">■ Setting up the network and storage for a cluster environment■ Verifying the DNS entries for the systems on which Exchange will be installed
“Establishing an MSCS cluster” on page 191	<ul style="list-style-type: none">■ Reviewing general guidelines to establish an MSCS cluster.
“Creating the MSDTC resource (Exchange 2003)” on page 192	<ul style="list-style-type: none">■ Creating the MSDTC resource for Exchange 2003.

Table 11-1 Tasks for installing and configuring SFW with MSCS and Exchange

Objective	Tasks
“Installing SFW with MSCS/Failover Cluster option” on page 194	<ul style="list-style-type: none"> ■ Verifying the driver signing options for Windows 2003 systems ■ Installing SFW (automatic installation) ■ Installing Cluster Option for Microsoft Cluster Service (MSCS) (manual option) ■ Restoring driver signing options for Windows 2003 systems
“Configuring SFW disk groups and volumes” on page 201	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes
“Preparing the forest and domain (Exchange 2003)” on page 208	<ul style="list-style-type: none"> ■ Setting up the forest and domain prior to an Exchange 2003 installation
“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 208	<ul style="list-style-type: none"> ■ Adding a Volume Manager Disk Group resource to the Cluster Group so that you can install the First Storage Group database files to an SFW dynamic volume
“Installing Exchange Server” on page 209	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Microsoft Exchange Server installation
“Creating an Exchange virtual server group (Exchange 2003)” on page 210	<ul style="list-style-type: none"> ■ Forming a cluster group that includes the IP Address, Network Name, Volume Manager Disk Group, and Exchange 2003 System Attendant (SA) resources
“Moving Exchange databases and logs to shared storage (Exchange 2003)” on page 221	<ul style="list-style-type: none"> ■ Moving the transaction log to a new volume ■ Moving the mailbox and public stores to the new volumes
“Moving Exchange databases and logs to shared storage (Exchange 2007)” on page 223	<ul style="list-style-type: none"> ■ Using the Exchange Management Console to change log file and system file locations ■ Using the Exchange Management Console to change the database file location
“Implementing a dynamic mirrored quorum resource” on page 225	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the Cluster Group ■ Changing the quorum resource to a dynamic mirrored quorum resource.

Table 11-1 Tasks for installing and configuring SFW with MSCS and Exchange

Objective	Tasks
“ Verifying the cluster configuration ” on page 228	■ Moving the online cluster group to the second node and back to the first node

Reviewing the requirements

Verify that the following requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation:

- “[Disk space requirements](#)” on page 187
- “[System requirements](#)” on page 187

Supported software

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW in a Microsoft cluster on Windows Server 2003:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2003 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 11-4 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2003 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2003 x64 Standard Edition or Enterprise Edition (SP2 required for all editions) ■ Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition (SP2 required for all editions)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 11-5 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

Observe the following system requirements:

- One CD-ROM drive accessible to the system on which you are installing MSCS.

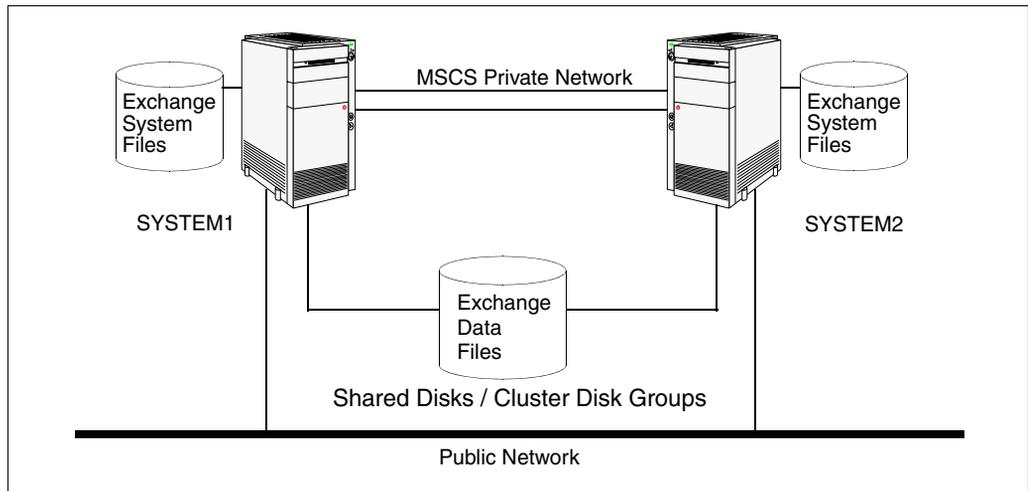
- Typical configurations require shared disks to support applications that migrate between nodes in the cluster. Symantec recommends two disks for Exchange: one for Exchange database files and one for Exchange log files.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Each system requires 1 GB of RAM for SFW.
- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- A minimum 256 MB of RAM per server is required for Exchange 2003; refer to your Microsoft documentation for more information.
- SFW requires administrator privileges to install the software.
- Before you install SFW, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List at the following URL to confirm supported hardware:
<http://www.symantec.com/business/support/index.jsp>
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft Exchange documentation for instructions on creating a reverse lookup zone.
- For Exchange 2003, you must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements during Microsoft procedures that do not involve Symantec wizards.

Reviewing the configuration

An example of a typical configuration for a cluster includes two servers and one storage array in an active/passive configuration. To set up one to one failover capabilities, EVS1 can fail over from SYSTEM1 to SYSTEM2 and vice versa.

[Figure 11-1](#) shows an example SFW configuration with Exchange and MSCS.

Figure 11-1 Example Exchange configuration



Some key points about the configuration:

- An MSCS cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.
 In an MSCS cluster without SFW, the quorum disk is a point of failure because MSCS only supports a basic physical disk and does not enable you to mirror the quorum resource.
 A key advantage of SFW is that it provides a dynamic mirrored quorum resource for MSCS. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables

you to verify that Exchange is working in the cluster before adding the dynamic quorum volume.

- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap.

See the *Veritas Storage Foundation Administrator's Guide*

Configuring the storage hardware and network

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.
- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before installing SFW, you must establish an MSCS cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To establish an MSCS cluster (general guidelines)

- 1 Verify that the quorum disk has been created before installing MSCS and configuring a cluster. (For IA64 systems, the quorum must be created using MBR instead of GPT or it will not be visible.)
- 2 Configure the shared storage and create a partition with drive letter "Q" for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster (SYSTEM1) using MSCS Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Verify that the node can access the shared storage.
- 4 Connect the second node to the shared storage.

- 5 Add the second node (SYSTEM2) using Cluster Administrator on that system.
- 6 Test the cluster by using the Move Group command to move the cluster resources to the second node.
SYSTEM2 becomes the active cluster node.

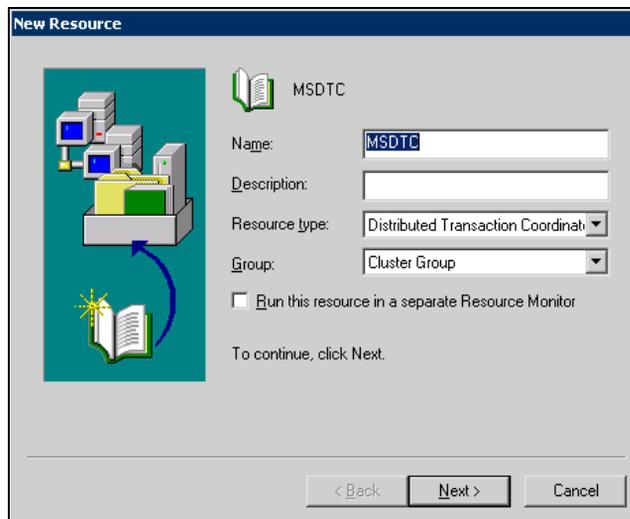
Creating the MSDTC resource (Exchange 2003)

This procedure is required for Exchange 2003.

Prior to installing Exchange 2003, create the MSDTC resource. You can create this resource now or just before installing Exchange.

To create the MSDTC resource

- 1 In Cluster Administrator (**Start > Administrative Tools > Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify the following options and then click **Next**:

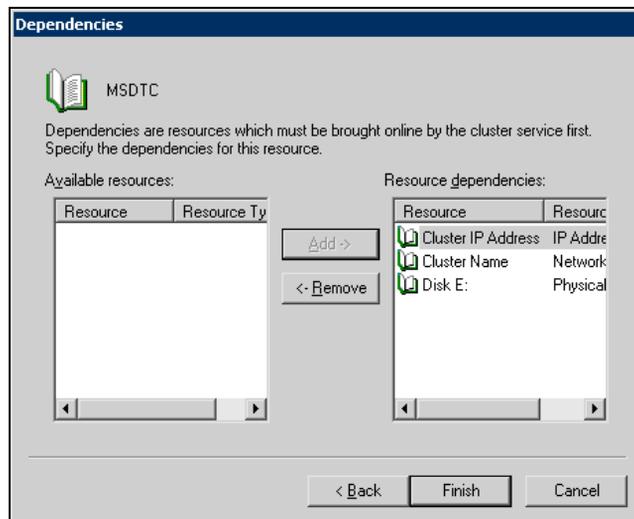


Name Type a name for the MSDTC resource.

Description If necessary, type a description about the resource

Resource type Click **Distributed Transaction Coordinator**.

- 3 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 4 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the Available Resources list, add them to the Resource dependencies list, and click **Finish**.



- 5 Click **OK**.
- 6 In the left pane, expand the Groups icon.
- 7 Click **Cluster Group**.
- 8 Right-click **Bring Online**.
The state changes to online.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running an MSCS cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the MSCS cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 194.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 196.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 199.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
See “[Changing the driver signing options](#)” on page 194.
- Moving the Online Groups
See “[Moving the online groups](#)” on page 196.

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The table below describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 11-6 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a MSCS configuration.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.
The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements. The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).
The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade

your software), contact the Symantec Technical Support, and see the contents of the DVD.

- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.
 - Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.
 If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.
 To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.
 Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** option.
 - Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
 - To add a computer for installation, provide the name of the computer in the System Name box.

OR

If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.

- To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
- To know the verification status and other information of the added computer, click the information icon.
- To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

■ **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

■ **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this

feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.
It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.
Click **Install** to install the software.
- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups
See [“Moving the online groups”](#) on page 199.
- Completing the SFW Installation
See [“Completing the SFW installation”](#) on page 200.
- Resetting the driver signing options
See [“Resetting the driver signing options”](#) on page 200.

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 194.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for Exchange. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different layouts.

Configuring disk groups and volumes involves the following tasks:

- “[Planning disk groups and volumes](#)” on page 201
- “[Creating a disk group](#)” on page 202
- “[Creating volumes](#)” on page 204

Planning disk groups and volumes

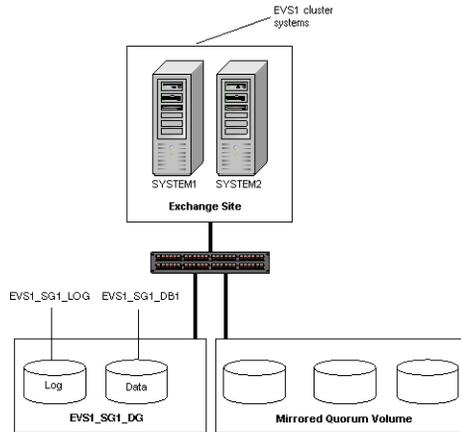
Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW. This is also an opportunity to grow existing volumes, add storage groups, and create volumes to support additional databases for existing storage groups.

Before you create a disk group, consider the following items:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs which depend on the traffic load.
- The number of disk groups and volumes that are needed for Exchange. Typically an SFW disk group corresponds to an Exchange storage group, with a separate volume for each database and for the transaction log.
- The disk groups and volumes for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application. You can wait until after the environment is configured to convert the basic disk quorum into a dynamic mirrored volume; this enables you to verify that Exchange is working in the cluster before adding the dynamic quorum volume.

[Figure 11-2](#) shows a detailed view of the disk groups and volumes for Exchange.

Figure 11-2 SFW disk groups and volumes for Exchange virtual server EVS1 in MSCS setup



Exchange storage group EVS1_SG1_DG contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups.

Creating a disk group

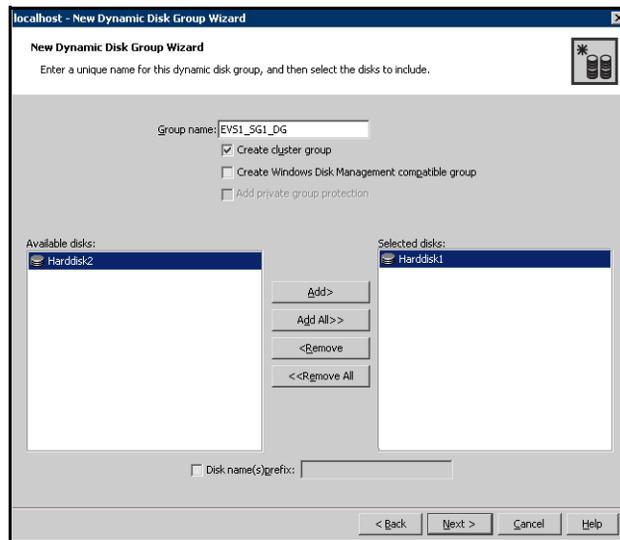
Use the following procedure to create a dynamic cluster disk group for an Exchange storage group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.

- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

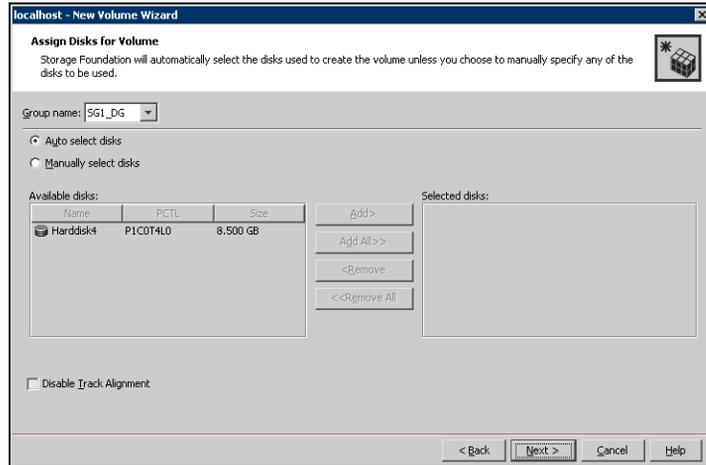
Creating volumes

Use the following procedure to create dynamic volumes. Typically you create a separate volume for each database and for the transaction log.

To create dynamic volumes

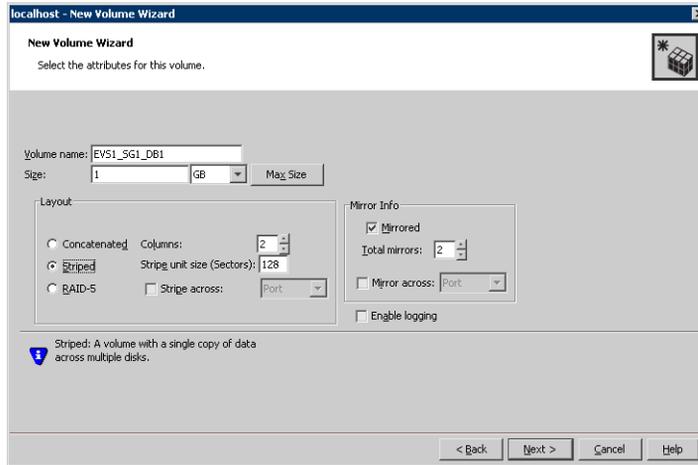
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example EVS1_SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.



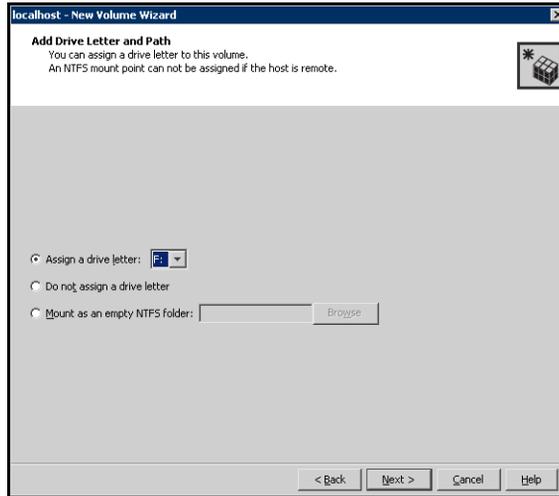
- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
 You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the volume attributes.



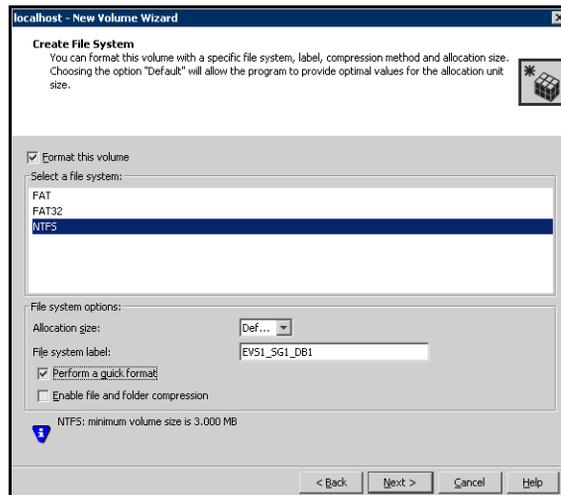
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.

- Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
- 14 Repeat these steps to create the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create any additional volumes required. Create the cluster disk group and volumes on the first node of the cluster only.

Preparing the forest and domain (Exchange 2003)

This procedure applies to Exchange 2003.

Microsoft requires the preparation of the forest and domain prior to an Exchange installation. See Microsoft documentation for instructions for preparing the forest and domain for an Exchange installation. Do not repeat this process for additional Exchange installations.

Adding a Volume Manager Disk Group resource for Exchange 2007 installation

This procedure applies to Exchange 2007.

Before Exchange 2007 installation you add a Volume Manager Disk Group resource to the Cluster Group so that you can install the First Storage Group database files to an SFW dynamic volume.

After Exchange installation you will move that resource from the Cluster Group to the Exchange group that is created during Exchange 2007 installation.

To add a Volume Manager Disk Group resource for Exchange 2007

- 1 Verify that the Cluster Group is online on the node.
- 2 Right-click on the Cluster Group and select **New > Resource**. The New Resource window appears.
- 3 On the New Resources window, do the following:
 - Specify a name for the disk group resource in the **Name** field.
 - If required, add a description about the resource in the **Description** field.

- Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.
 - If necessary, use the drop-down list to select the appropriate MSCS group; the group should already be selected.
 - Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
 - Click **Next**.
- 4 On the Possible Owners screen, all the nodes in the cluster are listed as possible owners by default. Click **Next**.
 - 5 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)
 - 6 Make sure the appropriate SFW cluster disk group is selected from the drop-down list for the resource, and click **Finish**.

Installing Exchange Server

Exchange 2003 requires service pack 2.

This is a standard Microsoft Exchange Server installation. Refer to the Microsoft documentation for details on this installation.

Complete the following tasks before installing Exchange Server:

- For Exchange 2003, prepare the forest and domain.
See [“Preparing the forest and domain \(Exchange 2003\)”](#) on page 208 for instructions.
- Verify that all systems on which Exchange Server will be installed have IIS installed. For Exchange 2003, you must install SMTP, NNTP, and WWW services on all systems, and if you install Exchange on Windows 2003, make sure to install ASP.NET service. For Exchange 2007, you must install WWW services.
- For Exchange 2003, make sure that the MSDTC resource exists.
See [“Creating the MSDTC resource \(Exchange 2003\)”](#) on page 192 for instructions.
- Make sure that the same drive letter and path are available on all nodes and have adequate space for the installation. For example, if you install the Exchange program files in C:\Program Files\ExchSrvr on one node, you must install the files in C:\Program Files\ExchSrvr on all the other nodes. Use a drive letter on the local disk of each node in the cluster.

Creating an Exchange virtual server group (Exchange 2003)

Creating an Exchange 2003 virtual server in an MSCS cluster involves forming a cluster group that includes the following resources:

- IP Address resource
- Network Name resource
- Volume Manager Disk Group resources for the SFW disk groups
- Exchange System Attendant (SA) resource; this resource adds the other required Exchange resources.

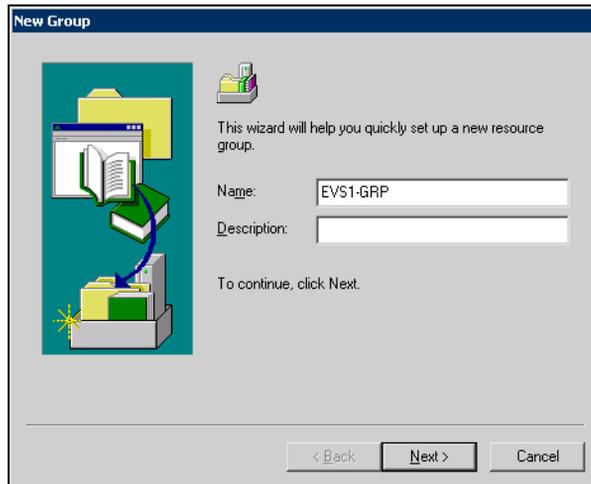
Note: During the installation of Exchange Server 2007, the IP Address resource, the Network Name resource, and the Exchange System Attendant (SA) resource are created. Therefore, you only need to add the Volume Manager Disk Group resources. See “[Adding Volume Manager Disk Group resources to the Exchange 2007 group](#)” on page 220.

Note: Before creating the resources, start the cluster service on all the nodes in the cluster.

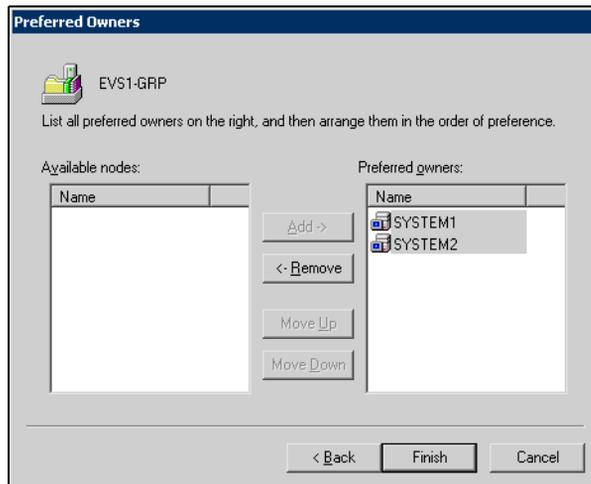
To configure the Exchange virtual server

- 1 In Cluster Administrator (**Start > Control Panel > Administrative Tools > Cluster Administrator**), connect to the appropriate cluster through the console.
- 2 In the configuration tree, right-click **Groups**, click **New**, and click **Group**.

- 3 In the New Group dialog box, specify a name for the group (EVS1-GRP), and click **Next**.



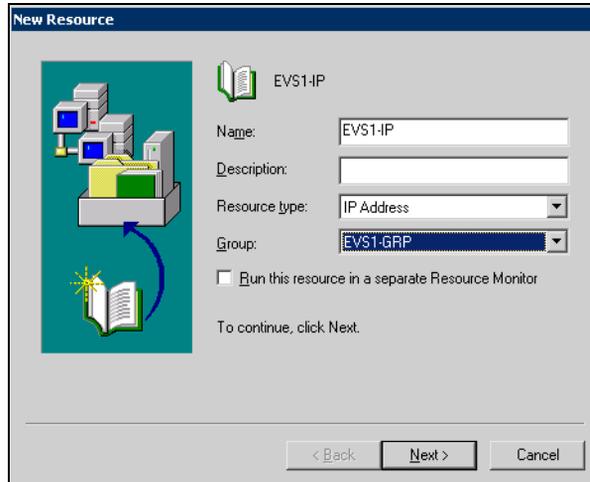
- 4 In the Preferred Owners dialog box, make sure that all the preferred owners are added to the Preferred Owners list, click **Finish** to create the group, and click **OK**.



- 5 Proceed to add resources to the group.

To create an IP Address resource for the Exchange virtual server

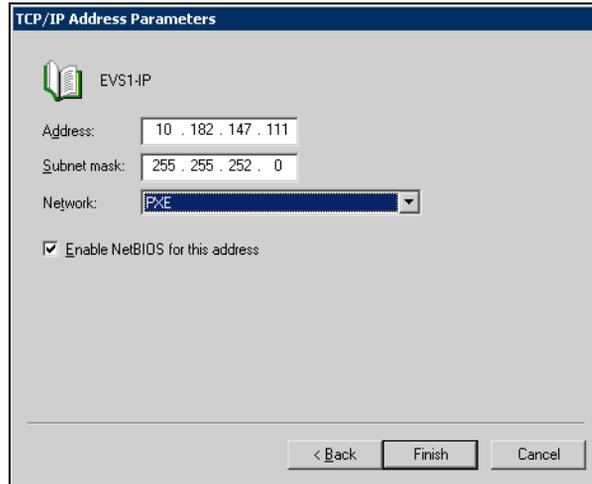
- 1 In the Cluster Administrator configuration tree, right-click the new Exchange virtual server group (EVS1-GRP), click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify the following information, and click **Next**.



Name	Type a name for the IP Address resource (EVS1-IP).
Description	If necessary, add a description about the resource.
Resource type	Click IP Address .

- 3 In the Possible Owners dialog box, click **Next**.
- 4 In the Dependencies dialog box, verify that the Resource dependencies list is empty, and click **Next**.

- 5 In the TCP/IP Address Parameters dialog box, enter the IP address and subnet mask, and then click **Finish**.

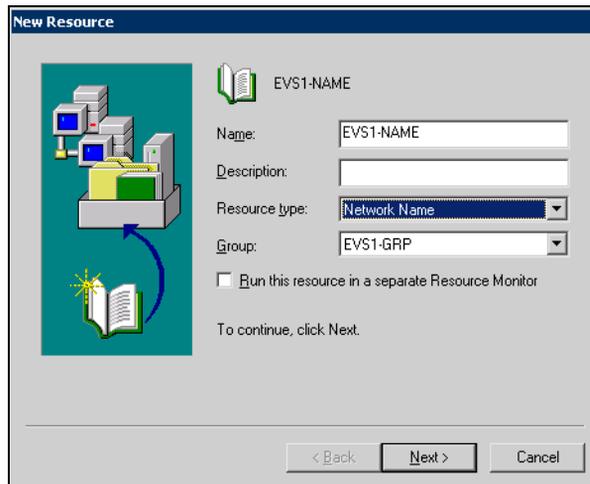


- 6 Click **OK**.

To create a Network Name resource for EVS

- 1 In the Cluster Administrator configuration tree, right-click the new Exchange virtual server group (EVS1-GRP), click **New**, and click **Resource**.

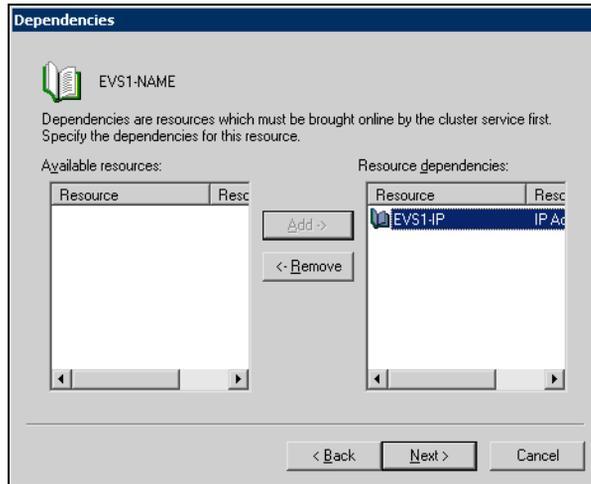
- 2 In the New Resource dialog box, specify the following, and click **Next**.



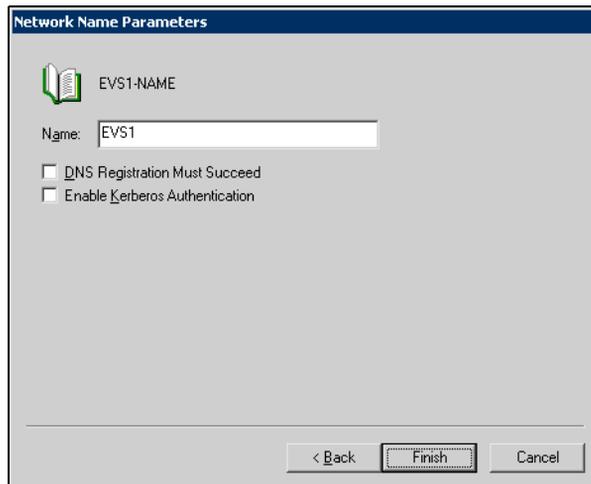
Name	Specify a name for the Network Name resource (EVS1-NAME).
Description	If necessary, type a description about the resource.
Resource type	Click Network Name .

- 3 In the Possible Owners dialog box, click **Next**.

- 4 In the Dependencies dialog box, select the IP Address resource from the Available resources list, add it to the Resource dependencies list, and click **Next**.



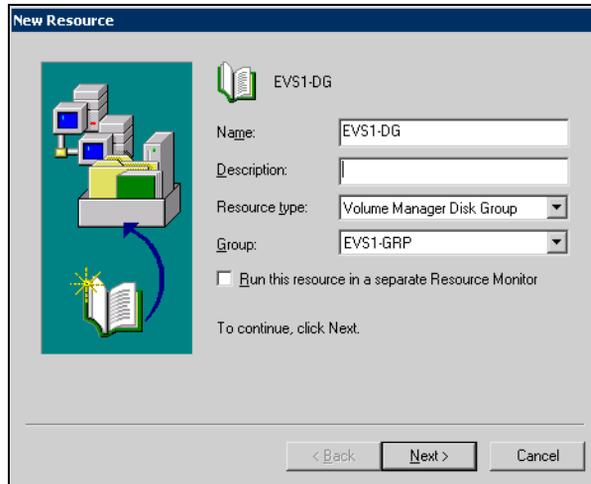
- 5 In the Network Name Parameters dialog box, enter the name of the Exchange virtual server, and click **Finish**.



- 6 Click **OK**.

To create a Disk Group resource

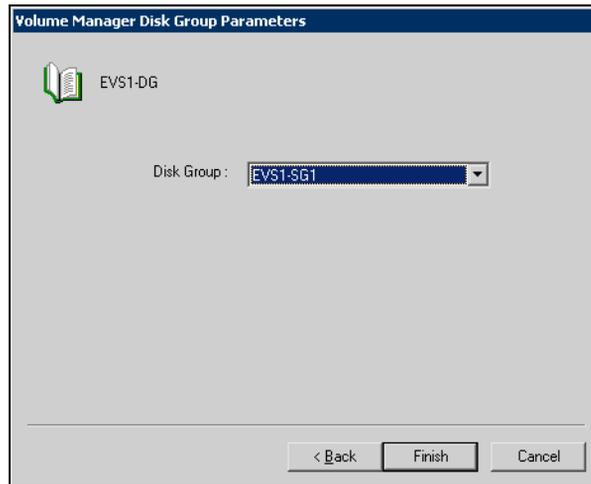
- 1 In the Cluster Administrator configuration tree, right-click the new Exchange virtual server group (EVS1-GRP), click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify the following, and then click **Next**.



Name	Specify a name for the Disk Group resource (EVS1-DG).
Description	If necessary, type a description about the resource.
Resource type	Click Volume Manager Disk Group .

- 3 In the Possible Owners dialog box, click **Next**.
- 4 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a Disk Group resource.

- 5 In the Volume Manager Disk Group Parameters dialog box, select the SFW disk group and click **Finish**.

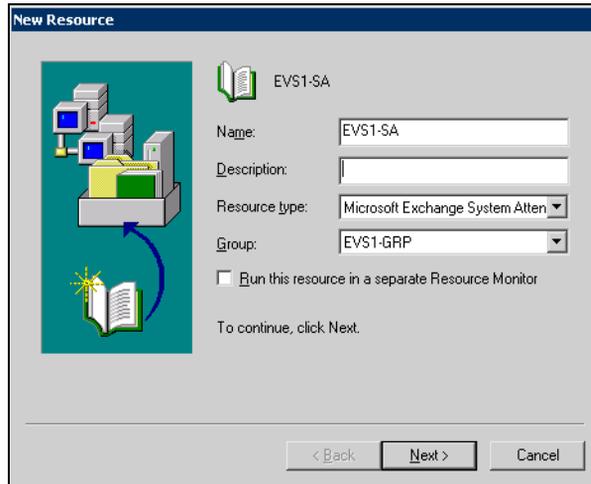


- 6 Click **OK**.

To create the Exchange 2003 SA resource

- 1 Before creating this resource, verify that the EVS IP resource, Network Name resource and the Disk Group resources are online. To bring the resources online, in the Cluster Administrator configuration tree, right-click the group (EVS1-GRP), and click **Bring Online**.
- 2 In the Cluster Administrator configuration tree, right-click the new Exchange virtual server group (EVS1-GRP), click **New**, and click **Resource**.

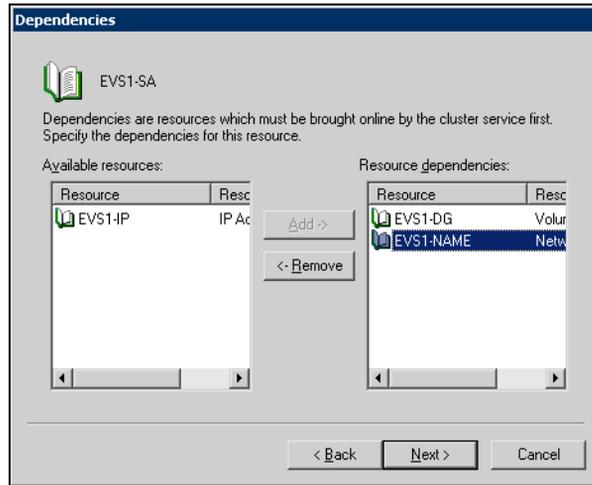
- 3 In the New Resource dialog box, specify the following, and then click **Next**.



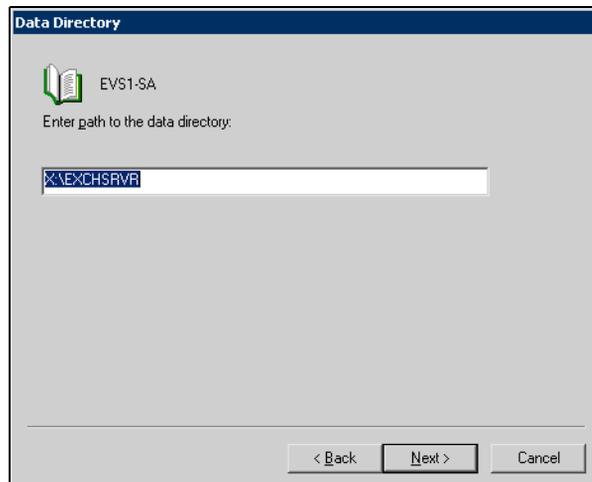
Name	Specify a name for the Exchange 2003 SA resource (EVS1-SA).
Description	If necessary, type a description about the resource.
Resource type	Click Microsoft Exchange System Attendant .

- 4 In the Possible Owners dialog box, click **Next**.

- 5 In the Dependencies dialog box, select the Network Name and Volume Manager Disk Group resources from the Available resources list, add them to the Resource dependencies list, and click **Next**.

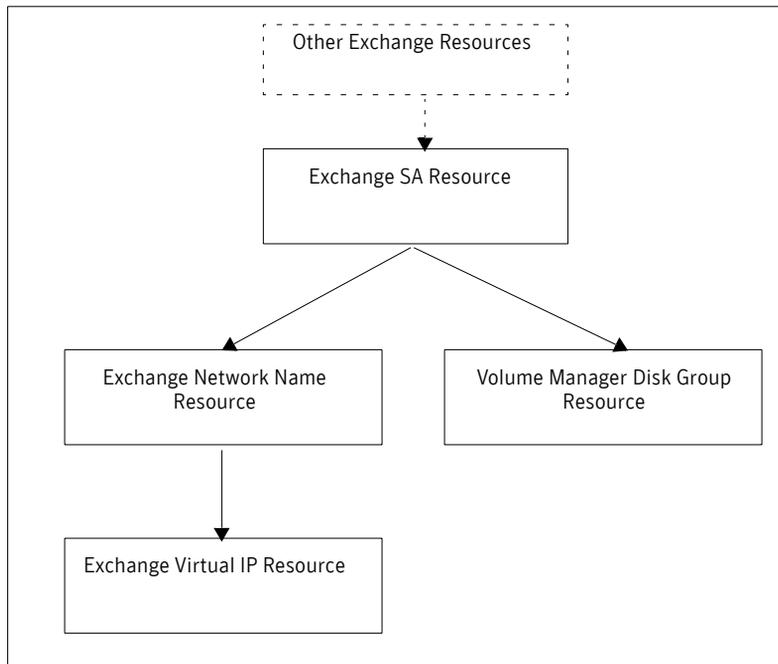


- 6 In the Exchange Administrative Group dialog box, click **Next**.
- 7 In the Exchange Routing Group dialog box, click **Next**.
- 8 In the Data Directory dialog box, verify that the drive letter displayed is the drive letter of the volume associated with the Exchange database (Exch_Data_SG1) in the SFW disk group (Exch_Group) that the Exchange 2003 SA resource is dependent on, and then click **Next**.



- 9 Click **Finish**.
- 10 Click **OK** after the Exchange 2003 SA resource is created.
- 11 Right-click the group (EVS1-GRP) in the Cluster Administrator configuration tree, and click **Bring Online** to bring the resources online. Creating the Exchange 2003 SA resource automatically adds the other required Exchange resources.
[Figure 11-3](#) shows an example dependency graph after creating the Exchange 2003 SA resource.

Figure 11-3 Dependency graph after creating the SA resource



Adding Volume Manager Disk Group resources to the Exchange 2007 group

For Exchange 2007 you add a Volume Manager Disk Group resource to the Cluster Group before installing Exchange. By adding the resource before installation, you can install the First Storage Group to a dynamic volume.

See [“Adding a Volume Manager Disk Group resource for Exchange 2007 installation”](#) on page 208.

Installing Exchange 2007 in the MSCS cluster automatically sets up a group for Exchange 2007 with the required resources for Exchange 2007.

Use the Cluster Administrator to move the Volume Manager Disk Group resource that you added for Exchange 2007 from the Cluster Group to the Exchange group. After doing so, set the Exchange database resource for the First Storage Group to depend on the Volume Manager Disk Group resource.

If you have created additional cluster disk groups for Exchange storage groups, add a Volume Manager Disk Group resource for each cluster disk group to the Exchange group. Set the appropriate dependencies.

Moving Exchange databases and logs to shared storage (Exchange 2003)

When Exchange Server is installed, the First Storage Group is created in the installation location, which by default is the boot drive. After you create the volumes in SFW, you must move the components (mailbox store and public folder store) of the First Storage Group to new volumes not on the boot drive.

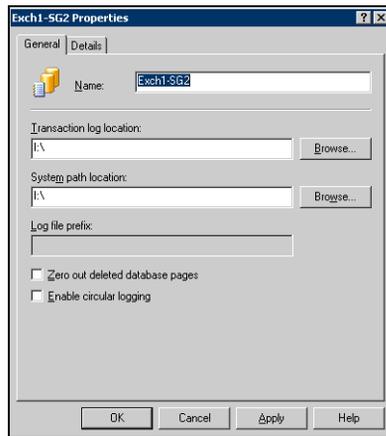
For Exchange 2003, set the path for the transaction log location and system path location fields to point to the log volume. Set the path for the Exchange database and Exchange streaming database files to point to the appropriate database volumes. You use the Exchange System Manager to specify the locations.

Optional steps for creating a new storage group and mailbox stores are included in the procedure.

To point the Exchange 2003 databases and log paths to the SFW volumes

- 1 Click **Start > All Programs > Microsoft Exchange > System Manager** to open the Exchange 2003 System Manager.
- 2 In the appropriate Administrative Group, expand **Servers** and expand the appropriate Exchange server.
- 3 Choose one of the following:
 - Right-click an existing storage group, for example SG2, and click **Properties**.
 - To create a new storage group, right-click the appropriate Exchange server, click **New** and click **Storage Group**.

4 Complete the following on the Properties sheet and click **OK**:

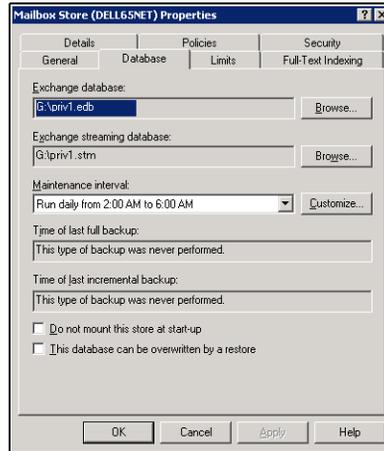


- | | |
|--------------------------|---|
| Name | For a new storage group, enter the name storage group name (for example, SG2). |
| Transaction log location | Click the Browse button and set the transaction log path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \. |
| System path location | Click the Browse button and set the path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \.

The paths for the Transaction log location and System path location should be the same. |
| Enable circular logging | Make sure that the Enable circular logging check box is not checked. |

- 5 Right-click on the storage group and choose one of the following:
- For an existing storage group, right-click on a mailbox store and click **Properties**.
 - For a new storage group, click **New** and click **Mailbox Store**.
- 6 Choose one of the following:
- For an existing storage group, proceed to [step 7](#).
 - For a new storage group, in the **General** tab of the Properties sheet enter the name of the new mailbox store (for example, SG2-DB1).

- 7 Click the **Database** tab, set the paths for the .edb and .stm files for the database as follows, and click **OK**:



Exchange database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing EDB files (for example, G: \).

Exchange streaming database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing STM files (for example, G: \).

- 8 Click **Yes** to mount the store.
- 9 Repeat [step 5](#) through [step 8](#) to create or set the paths for other mailbox stores. For example, create another mailbox store mounted on the H: drive, SG2-DB2.

Moving Exchange databases and logs to shared storage (Exchange 2007)

During Exchange 2007 installation, the First Storage Group database and log are installed on the same volume on shared storage. You move them to separate volumes after installation.

For each Exchange 2007 storage group, set the log files path and system files path to point to the log volume. For each database, set the database path to point to the appropriate database volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new storage groups or databases or change the locations for existing storage groups or databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

The following procedures can be used to change paths for existing storage groups and databases.

See the Microsoft Exchange 2007 documentation for additional information on creating new storage groups and databases.

To use the Exchange Management Console to change log file and system file locations for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the storage group for which you want to change the log file location.
- 4 In the work pane, right-click the storage group for which you want to change the log file location and click **Move Storage Group Path**. The Move Storage Group Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the log files, and create a folder for the log files. Repeat for the system files. The volume should be the same for both.
- 6 Click **Move**. A warning appears that all databases in the storage group must be temporarily dismounted, which will make them inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Storage Group Path wizard.

To use the Exchange Management Console to change the database file location for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the database for which you want to change the database file location.
- 4 In the work pane, right-click the desired database and click **Move Database Path**. The Move Database Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the database file, and create a folder for it as needed.
- 6 Click **Move**. A warning appears indicating that the database must be temporarily dismounted, which will make it inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the database files were moved successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Implementing a dynamic mirrored quorum resource

One of the key advantages of using SFW with MSCS is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster.

Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Create the quorum resource for the Cluster Group
- Change the quorum resource to a dynamic mirrored quorum resource

Note: If you are using DMP, you must create a dynamic quorum resource in order for the groups to failover properly.

Creating a dynamic cluster disk group for the quorum resource with mirrored volume

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a three-way mirrored volume in the New Volume wizard, select the Concatenated layout, select the Mirrored check box, and specify three mirrors. See “[Configuring SFW disk groups and volumes](#)” on page 201 for details on a creating cluster disk groups and volumes.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

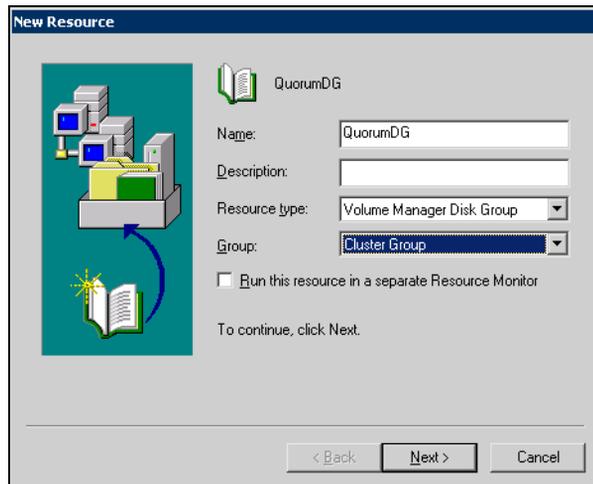
Creating the quorum resource for the cluster group

You must create the quorum resource for the cluster group.

To create the quorum resource for the cluster group

- 1 From Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**), verify that the Cluster Group is online on the same node where you created the disk group.
- 2 Right-click the **Cluster Group**, click **New**, and click **Resource**.

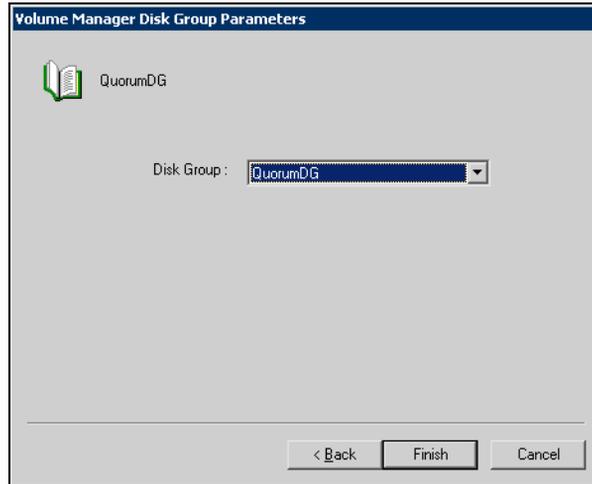
- 3 In the New Resource dialog box, specify the following, and then click **Next**.



Name	Specify a name for the quorum resource (QuorumDG).
Description	If necessary, type a description about the resource.
Resource type	Click Volume Manager Disk Group .

- 4 In the Possible Owners dialog box, click **Next**.
- 5 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.

- 6 In the Volume Manager Disk Group Parameters dialog box, select the disk group, and then click **Finish**.



- 7 Click **OK**.
- 8 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

Change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum resource to a dynamic mirrored quorum resource

- 1 In Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 In the Properties window, select the Quorum tab.
- 3 Select the name of the dynamic quorum disk group resource added in [step 2](#) on page 226.
- 4 Click **OK**.

Verifying the cluster configuration

To complete the configuration, verify that failover occurs normally in the cluster.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node. Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Deploying SFW with Microsoft failover clustering: New Exchange installation

This chapter includes the following topics:

- [Tasks for a new Exchange installation with SFW and failover clustering \(Windows Server 2008\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Establishing a Microsoft failover cluster](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Configuring SFW disk groups and volumes](#)
- [Implementing a dynamic mirrored quorum resource](#)
- [Adding a Volume Manager Disk Group resource for Exchange 2007 installation](#)
- [Installing Exchange Server](#)
- [Adding the Volume Manager Disk Group resources to the Exchange group](#)
- [Setting the database dependency on the disk group resource](#)
- [Moving Exchange databases and logs to shared storage](#)
- [Verifying the cluster configuration](#)

Tasks for a new Exchange installation with SFW and failover clustering (Windows Server 2008)

You can install and configure Veritas Storage Foundation for Windows (SFW) with Microsoft failover clustering and Microsoft Exchange Server 2007 on Windows Server 2008. The environment involves an active/passive configuration with one to one failover capabilities.

Exchange 2007 also runs on Windows Server 2003. For information on deploying SFW with Exchange Server 2007 on an MSCS cluster on Windows Server 2003, see:

- [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#) on page 183

[Table 12-1](#) outlines the high-level objectives and accompanying tasks that you perform, in the typical order in which you perform them.

Table 12-1 Tasks for deploying SFW and Exchange with failover clustering

Objective	Tasks
“Reviewing the requirements” on page 233	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 236	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster ■ Reviewing the benefits of a dynamic mirrored quorum
“Configuring the storage hardware and network” on page 238	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“Establishing a Microsoft failover cluster” on page 239	<ul style="list-style-type: none"> ■ Establishing a Microsoft failover cluster.
“Installing SFW with MSCS/Failover Cluster option” on page 242	<ul style="list-style-type: none"> ■ Installing SFW (automatic installation) ■ Installing Cluster Option for Microsoft Failover Clustering (manual option)
“Configuring SFW disk groups and volumes” on page 248	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes

Table 12-1 Tasks for deploying SFW and Exchange with failover clustering

Objective	Tasks
“Implementing a dynamic mirrored quorum resource” on page 256	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Adding the Volume Manager Disk Group resource for the quorum disk group ■ Changing the quorum resource to a dynamic mirrored quorum resource.
“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 259	<ul style="list-style-type: none"> ■ Adding a Volume Manager Disk Group resource for the SFW disk group that was created for the First Storage Group. You add it to the quorum group and then later move it to the Exchange group. This allows the First Storage Group to be installed to a dynamic disk.
“Installing Exchange Server” on page 259	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Microsoft Exchange Server installation ■ Installing the Active Clustered Mailbox role on the first node and the Passive Clustered Mailbox role on the failover node according to Microsoft instructions
“Adding the Volume Manager Disk Group resources to the Exchange group” on page 260	<ul style="list-style-type: none"> ■ Moving the Volume Manager Disk Group resource for the First Storage Group to the Exchange Group ■ Adding any other Volume Manager Disk Group resources needed for the Exchange databases
“Setting the database dependency on the disk group resource” on page 260	<ul style="list-style-type: none"> ■ Setting the database dependency on the Volume Manager Disk Group resource
“Moving Exchange databases and logs to shared storage” on page 261	<ul style="list-style-type: none"> ■ As necessary, altering the paths for the transaction log and databases
“Verifying the cluster configuration” on page 263	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node

Reviewing the requirements

Verify that the following requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation:

- [“Supported software for Microsoft failover clusters with SFW”](#) on page 234

- “[Disk space requirements](#)” on page 235
- “[System requirements](#)” on page 236

Supported software for Microsoft failover clusters with SFW

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW and Microsoft clustering on Windows Server 2008:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2008 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 12-2 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Table 12-2 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2008 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 12-3 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

Observe the following system requirements:

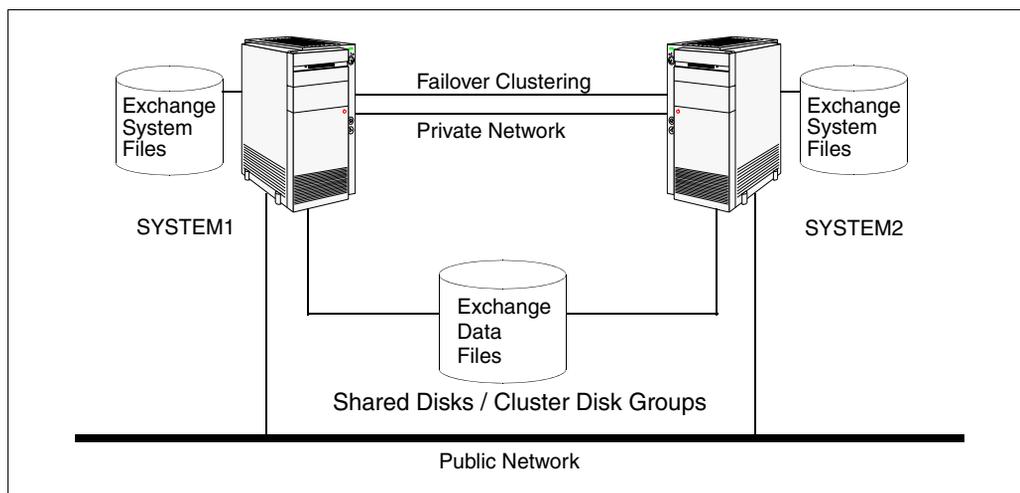
- One CD-ROM drive accessible to the system on which you are installing Microsoft failover clustering.
- Typical configurations require shared disks to support applications that migrate between nodes in the cluster. Symantec recommends two disks for Exchange: one for Exchange database files and one for Exchange log files.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- Microsoft failover clustering requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Each system requires 1 GB of RAM for SFW.
- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- SFW requires administrator privileges to install the software.
- Before you install SFW, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List at the following URL to confirm supported hardware:
<http://www.symantec.com/business/support/index.jsp>
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft Exchange documentation for instructions on creating a reverse lookup zone.

Reviewing the configuration

An example of a typical configuration for a cluster includes two servers and one storage array in an active/passive configuration. To set up one to one failover capabilities, the Exchange clustered mailbox server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

[Figure 12-1](#) shows an example SFW configuration with Exchange and Microsoft failover clustering.

Figure 12-1 Example Exchange configuration with Microsoft failover clustering



Some key points about the configuration:

- A Microsoft failover cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.

In a failover cluster without SFW, the quorum disk is a point of failure because failover clustering only supports a basic physical disk and does not enable you to mirror the quorum resource.

A key advantage of SFW is that it provides a dynamic mirrored quorum resource for failover clustering. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose.

You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables

you to verify that Exchange is working in the cluster before adding the dynamic quorum volume.

- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap.

See the *Veritas Storage Foundation Administrator's Guide*.

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.

- In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.

- 13 Close the window.

Establishing a Microsoft failover cluster

Before installing SFW, you must first verify that Microsoft failover clustering is enabled (if a new installation of Windows Server 2008), and then establish a Microsoft failover cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To enable Microsoft failover clustering

- 1 In Server Manager, select **Features** in the left pane (tree view) and then click **Add Features** (link on the right side of the screen).

- 2 In the Add Features Wizard, check the **Failover Clustering** option, and click **Next**.
- 3 Click **Install**.
- 4 When the installation is complete, click **Close**.

To establish a Microsoft failover cluster

- 1 Ensure that you have met the hardware prerequisites for a failover cluster. You can run the Microsoft wizard to validate the configuration. See the Microsoft documentation for details.
- 2 Configure the shared storage and create a volume with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends a minimum of 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster using Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 4 In the action pane, click **Create a Cluster**. The Create Cluster Wizard will start.
If this is the first time this wizard has been run, the Before You Begin page will appear. Review the information that is displayed and then click **Next**. You can hide this page on subsequent uses of the wizard; if this has been done, the first page might be the Select Servers page.
- 5 In the Select Servers panel, type the name of the first node in the Enter server name field and click **Add**. You can also use the Browse button to browse the Active Directory for the computers you want to add. Repeat this step for the second node.
- 6 After both nodes have been added to the list of Selected Servers, click **Next**.
- 7 Based on the information on the validation warning screen, assess your hardware configuration, and select one of the options. Symantec recommends that you select **Yes**, which starts the Validate a Configuration wizard. Follow the wizard instructions.
- 8 In the Access Point for Administering the Cluster screen, in the Cluster Name field, type the name for the failover cluster. This is the name that you use to connect to and administer the cluster.
- 9 In the Address field of the network area, type the appropriate IP address and then click **Next**.

- 10 In the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.
- 11 Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the 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, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running a Microsoft failover cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 242.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 243.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 246.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups
See “[Moving the online groups](#)” on page 242.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.

If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a Microsoft failover cluster configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

Note: Before you install Storage Foundation for Windows, make sure that the node is inactive.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.
The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements.

The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).

The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.

- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.
 - Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.

If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.

To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.

Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** option.

- Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
- To add a computer for installation, provide the name of the computer in the System Name box.
OR
If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.
 - To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
 - To know the verification status and other information of the added computer, click the information icon.
 - To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

■ **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

■ **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.
It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.
Click **Install** to install the software.
- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
See "[Moving the online groups](#)" on page 247.
- Completing the SFW installation
See "[Completing the SFW installation](#)" on page 247.

Moving the online groups

You can move the resource groups from the current system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the Microsoft failover cluster.

See “[SFW installation tasks](#)” on page 242.

Configuring SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and volumes for Exchange. A dynamic disk group is a collection of one or more disks that behaves as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts.

Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic cluster disk groups](#)
- [Creating volumes](#)
- [Managing disk groups and volumes](#)

Planning disk groups and volumes

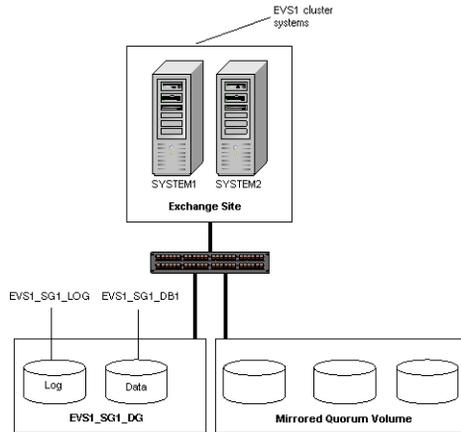
Before installing Exchange, you must create disk groups and volumes using the VEA console installed with SFW.

Before creating a disk group, consider:

- The type of volume configurations that are required.
- The number of LUNs required for the disk group.
- The implications of backup and restore operations on the disk group setup.
- The size of databases and logs that depend on the traffic load.
- The number of disk groups and volumes that are needed for Exchange. Typically an SFW disk group corresponds to an Exchange storage group, with a separate volume for each database and for the transaction log.
- The disk groups and volumes for the mirrored quorum resource
You will need a disk group with three disks for the mirrored quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. You can create the quorum disk group at the same time you create application disk groups, although it is not required for installing the application. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that SQL is working in the cluster before adding the dynamic quorum volume.

[Figure 12-2](#) shows a detailed view of the disk groups and volumes for Exchange.

Figure 12-2 SFW disk groups and volumes for Exchange virtual server (clustered mailbox server) EVS1 in Microsoft clustering setup



Exchange storage group EVS1_SG1_DG contains two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

The general guidelines for disk group and volume setup for EVS1_SG1_DG also apply to additional storage groups.

Creating dynamic cluster disk groups

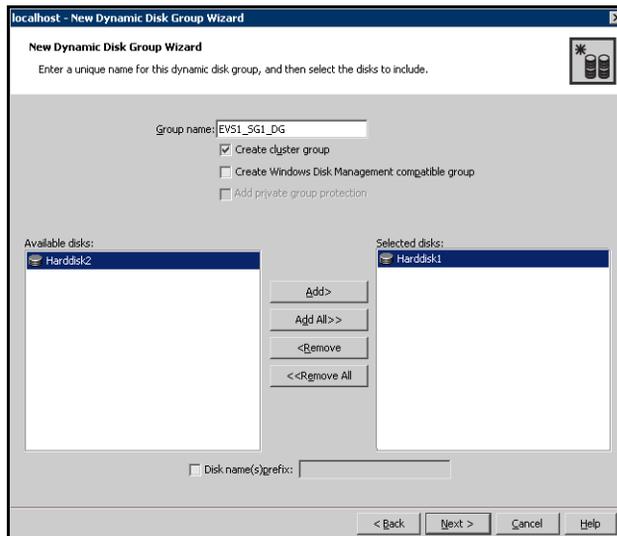
Use the following procedure to create a dynamic cluster disk group for an Exchange storage group.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

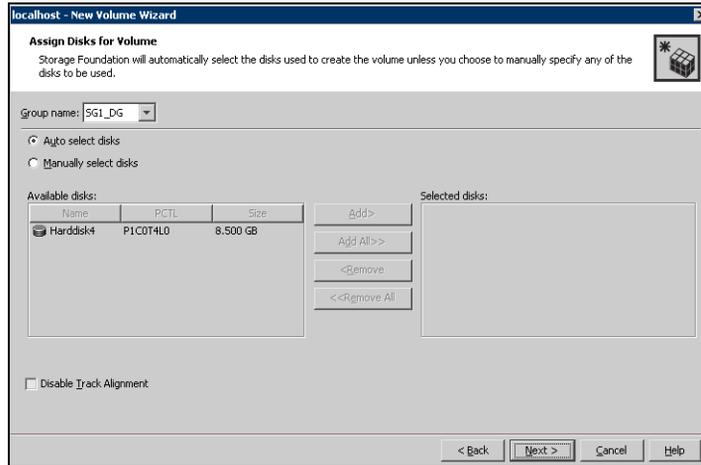
Creating volumes

Use the following procedure to create dynamic volumes. Typically you create a separate volume for each database and for the transaction log.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example EVS1_SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

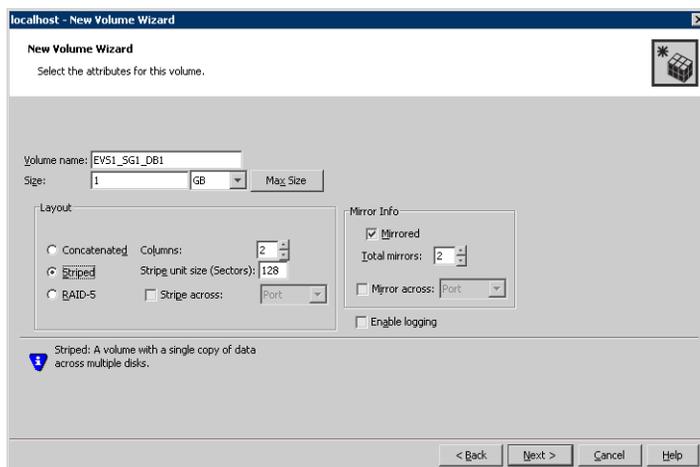


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

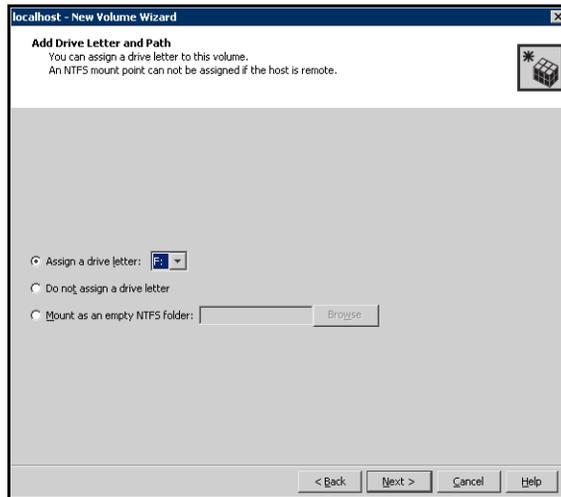
- 8 Click **Next**.

9 Specify the volume attributes.



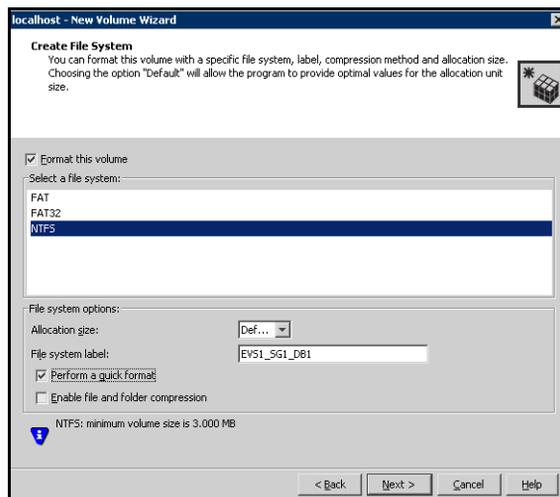
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.

- Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
- 14 Repeat these steps to create the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) Also, repeat these steps to create any additional volumes required. Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.

- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter
Select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Implementing a dynamic mirrored quorum resource

One of the key advantages of using SFW with Microsoft clustering is the ability to create a mirrored quorum resource that adds fault tolerance to the quorum and protects the cluster.

Complete the following tasks:

- Create a dynamic cluster disk group for the quorum resource with a mirrored volume
- Add the Volume Manager Disk Group resource for the quorum
- Change the quorum resource to a dynamic mirrored quorum resource

Note: If you are using DMP, you must create a dynamic quorum resource in order for the groups to failover properly.

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using three (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a three-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with three small disks.
- 2 Create a volume with the three disks.
- 3 Select the **Concatenated** layout, click the **Mirrored** check box, and specify three mirrors.

For full details on creating cluster disk groups and volumes, see [“Configuring SFW disk groups and volumes”](#) on page 248.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by Microsoft clustering.

Adding the Volume Manager Disk Group resource for the quorum

You must add the Volume Manager Disk Group resource corresponding to the disk group that you created for the quorum.

To add a Volume Manager Disk Group resource for the quorum in a Windows Server 2008 cluster

- 1 If Failover Cluster Management is already open, then proceed to Step 2. To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.
- 2 Verify that the cluster is online on the same node where you created the disk group.
- 3 In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
- 4 Right-click the new group and rename it, for example QUORUM.
- 5 Right-click QUORUM and select **Add a resource > More resources > Add Volume Manager Disk Group**.

- 6 Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
- 7 In the General tab of the Properties dialog box, type a name for the resource in the Resource Name field, for example, QUORUM_DG_RES.
- 8 On the Properties tab, in the Disk Group Name field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
- 9 Right-click the Quorum disk group resource (for example, QUORUM_DG_RES) in the left pane and select **Bring this resource online**.
The specified disk group resource, QUORUM_DG_RES resource, is created under the Quorum group (for example, QUORUM).

Changing the quorum resource to a dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

- 1 In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.
The Configure Cluster Quorum Wizard opens.
- 2 Review the screen and click **Next**.
- 3 Select either the **Node and Disk Majority** or **No Majority: Disk Only** radio button, and click **Next**.
- 4 Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.
This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, QUORUM_DG_RES.
- 5 Review the information in the Confirmation screen and click **Next**.
- 6 Click **Finish** to close the wizard.

Adding a Volume Manager Disk Group resource for Exchange 2007 installation

Before installing Exchange 2007, you add a Volume Manager Disk Group (VMDG) resource for the disk group that you created for the First Storage Group. By doing so, you can install the First Storage Group on a dynamic volume.

You add this resource to the Quorum group temporarily. After installation, you will move it to the Exchange group created by installation and set the appropriate dependencies.

Before creating this resource, start the cluster service on all the nodes in the cluster.

To create a Volume Manager Disk Group resource for the application

- 1 In the left pane (tree view) of the Failover Cluster Management tool, right-click the Quorum resource group under Services and Applications. Select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 2 In the General tab of the Properties dialog box, type a resource name for the new Volume Manager Disk Group (for example, EVS1_SG1_RES).
- 3 In the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the First Storage Group (for example, EVS1_SG1_DG).
- 4 Click **OK** to close the dialog box.
- 5 To bring the resource online, right-click the newly named resource and click **Bring this resource online**.

Installing Exchange Server

Exchange 2007 requires service pack 1 for Windows Server 2008.

Refer to the Microsoft documentation for prerequisites and details on installing a clustered mailbox server. You must install the Active Clustered Mailbox role on the first node and install the Passive Clustered Mailbox role on the failover node.

Make sure that the same drive letter and path are available on all nodes and have adequate space for the installation. For example, if you install the Exchange program files in `C:\Program Files\ExchSrvr` on one node, you must install the files in `C:\Program Files\ExchSrvr` on all the other nodes. Use a drive letter on the local disk of each node in the cluster.

Adding the Volume Manager Disk Group resources to the Exchange group

Exchange 2007 installation creates an application group with resources required by Exchange.

You need to move to the Exchange Group the Volume Manager Disk Group (VMDG) resource that you added to the cluster quorum group for the First Storage Group installation.

In addition, add to the Exchange group the Volume Manager Disk Group resources for any other disk groups that you created for Exchange.

To move the Volume Manager Disk Group resource to the application group

- 1 In the quorum group, right click the resource you created for the Exchange First Storage Group and select the option to move it to another group.
- 2 Select the Exchange group as the target for the move and click **OK**.

To add Volume Manager Disk Group resources for the application

- 1 In the left pane (tree view) of the Failover Cluster Management tool, right-click the resource group under Services and Applications. Select **Add a resource > More resources > Add Volume Manager Disk Group**. A Volume Manager disk group resource is automatically created.
- 2 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** (in the center pane) to open the Properties dialog box. You can also right-click **New Volume Manager Disk Group** and select **Properties**.
- 3 In the General tab of the Properties dialog box, type the resource name for the New Volume Manager Disk Group (for example, ExchDG).
- 4 In the Properties tab, type the name of the disk group for which you want to add a resource.
- 5 Click **OK** to close the dialog box.
- 6 To bring the resource online, right-click the newly named disk group and click **Bring this resource online**.

Setting the database dependency on the disk group resource

After adding the Volume Manager Disk Group resources to the Exchange group, you must set the Exchange database resources to be dependent on them.

To set the database resource dependency on the VMDG resource

- 1 In Failover Cluster Management, select the Exchange resource group.
- 2 In the result pane, under Other Resources, right-click the appropriate database resource and select Properties.
- 3 In the Dependencies tab of the Properties dialog box:
 - Click the box **Click here to add a dependency**.
 - Select the appropriate Volume Manager Disk Group resource from the dropdown list of available resources.
 - Click **Insert**.
- 4 Click **OK** to close the Properties dialog box.
- 5 Repeat steps 2 through 4 for each additional database resource that exists in the Exchange group.

Moving Exchange databases and logs to shared storage

During Exchange installation, the First Storage Group is installed to a dynamic volume. You must move the log to a separate volume.

If you created any other Exchange storage groups that were not located on the SFW dynamic volumes, move them to the dynamic volumes.

For each Exchange 2007 storage group, set the log files path and system files path to point to the log volume. For each database, set the database path to point to the appropriate database volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new storage groups or databases or change the locations for existing storage groups or databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

The following procedures can be used to change paths for existing storage groups and databases.

See the Microsoft Exchange 2007 documentation for additional information on creating new storage groups and databases.

To use the Exchange Management Console to change log file and system file locations for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the storage group for which you want to change the log file location.
- 4 In the work pane, right-click the storage group for which you want to change the log file location and click **Move Storage Group Path**. The Move Storage Group Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the log files, and create a folder for the log files. Repeat for the system files. The volume should be the same for both.
- 6 Click **Move**. A warning appears that all databases in the storage group must be temporarily dismounted, which will make them inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Storage Group Path wizard.

To use the Exchange Management Console to change the database file location for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the database for which you want to change the database file location.
- 4 In the work pane, right-click the desired database and click **Move Database Path**. The Move Database Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the database file, and create a folder for it as needed.

- 6 Click **Move**. A warning appears indicating that the database must be temporarily dismantled, which will make it inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the database files were moved successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Verifying the cluster configuration

To complete the configuration, verify that failover occurs normally in the cluster.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Failover Cluster Management to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node. Do not simulate failover in a production environment.

To move online groups

- 1 Open the Failover Cluster Management tool (**Start > All Programs > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.

Verifying the cluster configuration

- 2 Open Failover Cluster Management. Click **Start > All Programs > Administrative Tools > Failover Cluster Management** from any node in the cluster.
- 3 In Failover Cluster Management, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move the resource groups back to the original node, restart the node you shut down in [step 1](#), select the resource group, and use **Move this service or application to another node > Move to node [name of node]** to move the resource group.

Deploying SFW with MSCS and Exchange in a campus cluster

This chapter discusses the following topics:

- [Tasks for a new Exchange, SFW and MSCS installation in a campus cluster \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Establishing an MSCS cluster](#)
- [Creating the MSDTC resource \(Exchange 2003 only\)](#)
- [Installing SFW](#)
- [Creating disk groups and volumes](#)
- [Implementing a dynamic quorum resource](#)
- [Adding a Volume Manager Disk Group resource for Exchange 2007 installation](#)
- [Installing the application on the cluster nodes](#)
- [Setting up a group for Exchange 2003 in MSCS](#)
- [Adding Volume Manager Disk Group resources to the Exchange 2007 group](#)
- [Moving Exchange databases and logs to dynamic volumes](#)
- [Verifying the cluster configuration](#)

Tasks for a new Exchange, SFW and MSCS installation in a campus cluster (Windows Server 2003)

On Windows Server 2003, you can install and configure Veritas Storage Foundation for Windows with MSCS and Microsoft Exchange Server 2003 or 2007. This chapter presents a campus clustering example using a two-node cluster.

For information on deploying SFW with Microsoft failover clustering, which runs under Windows Server 2008, see:

- [Chapter 14, “Deploying SFW with Microsoft failover clustering and Exchange in a campus cluster” on page 317](#)

[Table 13-1](#) lists the high-level objectives for deploying SFW with MSCS in a campus cluster for Exchange, as well as the tasks within each objective:

Table 13-1 Task list for deploying SFW with MSCS in a campus cluster

Objectives	Tasks
“Reviewing the requirements” on page 269	<ul style="list-style-type: none">■ Verify hardware and software prerequisites.
“Reviewing the configuration” on page 272	<ul style="list-style-type: none">■ Review the configuration requirements.■ Review the overview of MSCS campus cluster, and recovery scenarios.
“Configuring the network and storage” on page 280	<ul style="list-style-type: none">■ Install the hardware for Site A. The server and storage array are connected to the SAN. Leave the cables for the NICs unconnected, and do not yet connect the switch to site B.■ Install the hardware in the same manner for Site B.

Table 13-1 Task list for deploying SFW with MSCS in a campus cluster

Objectives	Tasks
“Establishing an MSCS cluster” on page 282	<ul style="list-style-type: none"> ■ Install and configure the operating system and MSCS on Server A. ■ Configure the storage and create a partition for the cluster quorum disk on Site A. ■ Create the first node of the cluster on Server A. ■ Install and configure the operating system and MSCS on Server B. ■ Connect the two nodes. ■ Create the second node of the cluster on Server B. ■ Test the cluster by moving the resources to Server B. Server B becomes the active node. Do not move them back to Server A at this point.
“Creating the MSDTC resource (Exchange 2003 only)” on page 285	<ul style="list-style-type: none"> ■ Create the MSDTC resource.
“Installing SFW” on page 286	<ul style="list-style-type: none"> ■ Install SFW on Node A (Node B active). ■ Install SFW on Node B (Node A active).
“Creating disk groups and volumes” on page 294	<ul style="list-style-type: none"> ■ In SFW on Node A, create two or more dynamic cluster disk groups on the storage, one or more for the application data files and one for the mirrored quorum.
“Implementing a dynamic quorum resource” on page 303	<ul style="list-style-type: none"> ■ If not done earlier, create a dynamic disk group for the quorum with a mirrored volume. ■ Make that disk group into a Volume Manager Disk Group type resource in the default Cluster Group. ■ Change the quorum resource to the dynamic mirrored quorum resource.
“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 306	<ul style="list-style-type: none"> ■ For Exchange 2007, add a VMDG resource that corresponds to the disk group you created for the First Storage Group. You can add this to the Cluster Group. Later you will move it to the Exchange group created by Exchange 2007 installation.

Table 13-1 Task list for deploying SFW with MSCS in a campus cluster

Objectives	Tasks
“Installing the application on the cluster nodes” on page 307	<ul style="list-style-type: none"> ■ Follow Microsoft instructions for installing Exchange on the cluster. ■ Install the application program files on the local drive. ■ For Exchange 2007, install files relating to the data and logs on the dynamic volume on shared storage.
“Setting up a group for Exchange 2003 in MSCS” on page 309	<ul style="list-style-type: none"> ■ Create a group within MSCS for the Exchange application. ■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group. ■ Include other resources required by the application. ■ Include the required dependencies. For example, the Exchange SA resource depends on the disk group resource.
“Adding Volume Manager Disk Group resources to the Exchange 2007 group” on page 310	<ul style="list-style-type: none"> ■ For Exchange 2007, move the Volume Manager Disk Group resource for the First Storage Group to the Exchange 2007 group that was created during Exchange installation. Add any other Volume Manager Disk Group resources for other disk groups as needed. Set each database dependency on the appropriate VMDG resource.
“Moving Exchange databases and logs to dynamic volumes” on page 310	<ul style="list-style-type: none"> ■ Move the First Storage Group database and log to dynamic volumes (Exchange 2003) or move the log to a separate dynamic volume (Exchange 2007).
“Verifying the cluster configuration” on page 314	<ul style="list-style-type: none"> ■ Verify the cluster configuration by switching service groups or shutting down an active cluster node

Reviewing the requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

See the following topics:

- [Supported software](#)
- [System requirements](#)
- [Disk space requirements](#)

Supported software

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW in a Microsoft cluster on Windows Server 2003:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2003 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 13-4 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2003 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2003 x64 Standard Edition or Enterprise Edition (SP2 required for all editions) ■ Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition (SP2 required for all editions)

System requirements

To deploy SFW with MSCS in a campus cluster, your system must meet the following requirements:

- One CD-ROM drive accessible to each system on which you are installing MSCS.
- SCSI or Fibre Channel host bus adapters (HBAs) to access the storage.
- MSCS requires at least two network adapters per system (one network adapter to connect each system to the public network and one network adapter for the private network on each system). Symantec recommends having two adapters for the private network and routing each private network adapter through a separate hub or switch to avoid single points of failure.
- Refer to application documentation to determine disk space requirements for your application.
- Each system requires 1 GB of RAM for SFW.

- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- A minimum 256 MB of RAM per server is required for Exchange 2003; refer to your Microsoft documentation for more information.
- The configuration requires two sites with a storage array for each site, with an equal number of disks at each site for the mirrored volumes.
- Interconnects between the clusters are required for the storage and the network.
- Systems to be clustered must be configured as part of a Windows Server 2003 domain. Each system in an MSCS cluster must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and MSCS software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. Six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.

Note: To determine the approved hardware for SFW, see the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp>.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 13-5 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

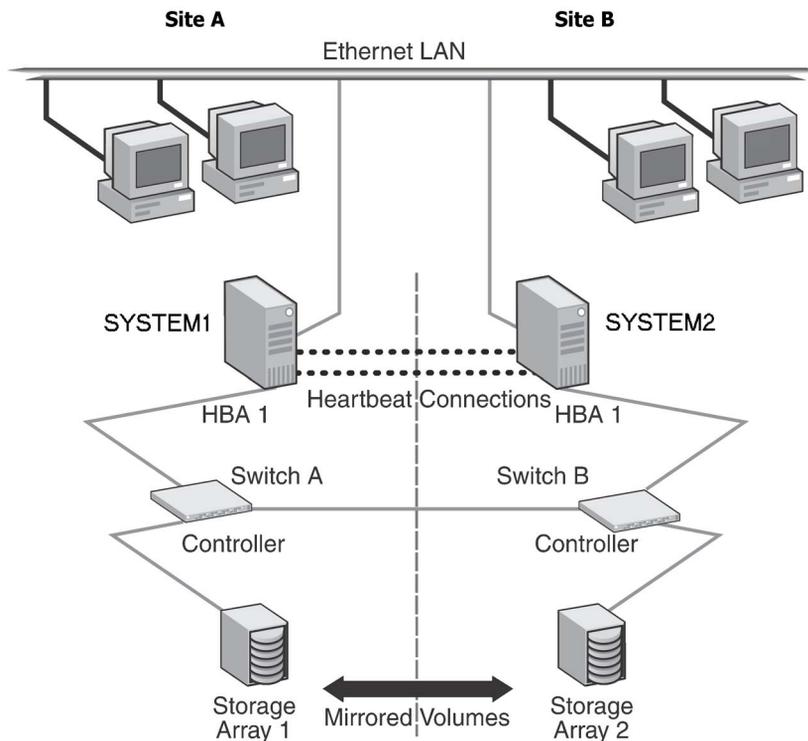
Note: Plan for an equal number of disks on the two sites, because each disk group should contain the same number of disks on each site.

Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

Figure 13-1 illustrates an example campus clustering configuration.

Figure 13-1 MSCS campus clustering configuration example



The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array and contains mirrored data of the storage on the other array. Each disk group should contain the same number of disks on each site for the mirrored volumes.

MSCS uses the quorum architecture, where the cluster database resides in the quorum resource. If you use MSCS for clustering, adding SFW to the configuration protects the quorum disk from being a single point of failure in the cluster because SFW provides dynamic volumes and software mirroring of the quorum device. To avoid a single point of failure, set up the quorum as a dynamic mirrored device. This example includes the dynamic mirrored quorum and requires setting up two or more dynamic cluster disk groups in SFW—one or more cluster disk groups for the application and data and one for the dynamic mirrored quorum.

The example configuration does not include Dynamic Multi-pathing (DMP). For instructions on how to add DMP to a clustering configuration, see *Veritas Storage Foundation and High Availability Solutions, Solutions Guide*.

When you are installing SFW and MSCS together, remember the following:

- To install SFW, an MSCS cluster must be running.
Before you install SFW, you must set up the hardware and install the operating system and MSCS on all systems and establish the MSCS cluster. Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Therefore, use a “rolling install” procedure to install SFW first on the inactive cluster node. Then move the cluster resources to the other node and install on the now inactive node.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

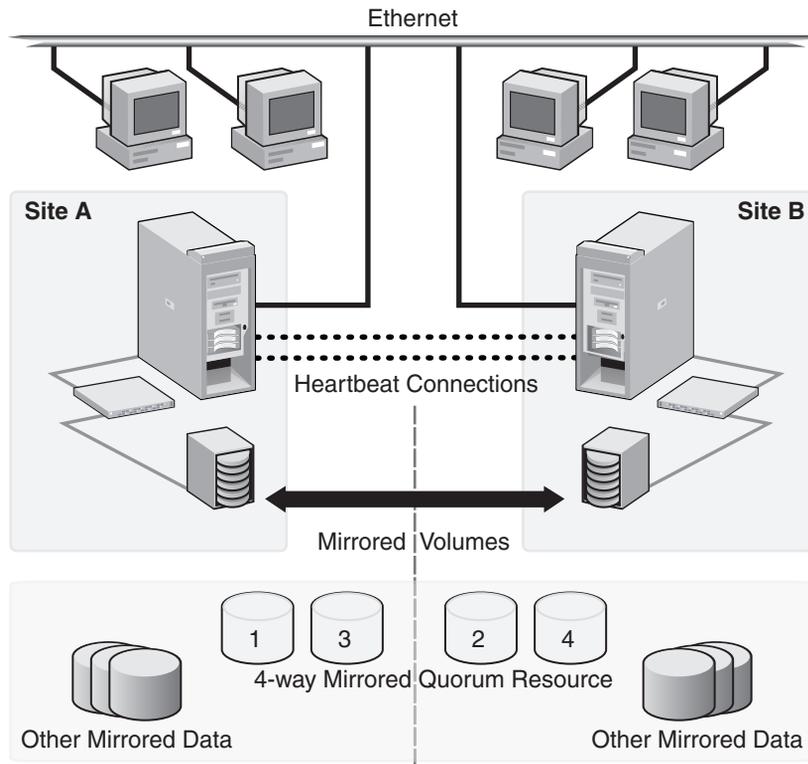
- Using SFW also offers other advantages over using MSCS alone. SFW lets you add fault tolerance to your data volumes. Mirroring of log volumes is recommended, and a mirrored striped RAID layout is recommended for your data volumes. SFW also offers multiple disk groups, multiple mirrors, capacity management and Automatic Volume Growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, Dynamic Multi-pathing, and enhanced snapshot capabilities with FlashSnap.

Overview of campus clustering with MSCS

[Figure 13-2](#) on page 274 shows an MSCS campus cluster configuration with mirrored storage across clusters and a mirrored quorum resource. The 4-way mirrored quorum has an extra set of mirrors for added redundancy. Although a

campus cluster setup with MSCS can work without Storage Foundation for Windows, SFW provides key advantages over using MSCS alone. Through a dynamic mirrored volume that functions on multiple disks across multiple sites, SFW protects the quorum resource in the cluster from being the single point of failure in the cluster.

Figure 13-2 Typical MSCS campus clustering configuration



Most customers use hardware RAID to protect the quorum disk, but that does not work when a natural disaster takes down the primary node and its attached storage. If the quorum resource is lost to the cluster, the cluster fails, because none of the cluster servers can gain control of the quorum resource and ultimately the cluster. MSCS alone cannot provide fault tolerance to the quorum disk.

MSCS campus cluster failure scenarios

This section focuses on the failure and recovery scenarios with an MSCS campus cluster and SFW installed.

For information about the quorum resource and arbitration in MSCS, see: [“MSCS quorum and quorum arbitration”](#) on page 278.

[Table 13-6](#) lists failure situations and the outcomes that occur:

Table 13-6 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Failover	If the services stop for an application failure, the application automatically fails over to the other site.
Server failure (Site A) May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	Failover	Assuming a two-node cluster pair, failing a single node results in a cluster failover. There will be a temporary service interruption for cluster resources that are moved from the failed node to the remaining live node.
Server failure (Site B) May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	No interruption of service.	Failure of the passive site (Site B) does not interrupt service to the active site (Site A).
Partial SAN network failure May mean that SAN fiber channel cables were disconnected to Site A or Site B Storage.	No interruption of service.	Assuming that each of the cluster nodes has some type of Dynamic Multi-pathing (DMP) solution, removing one SAN fiber cable from a single cluster node should not effect any cluster resources running on that node, because the underlying DMP solution should seamlessly handle the SAN fiber path failover.

Table 13-6 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p>Private IP Heartbeat Network Failure</p> <p>May mean that the private NICs or the connecting network cables failed.</p>	<p>No interruption of service.</p>	<p>With the standard two-NIC configuration for a cluster node, one NIC for the public cluster network and one NIC for the private heartbeat network, disabling the NIC for the private heartbeat network should not effect the cluster software and the cluster resources, because the cluster software will simply route the heartbeat packets through the public network.</p>
<p>Public IP Network Failure</p> <p>May mean that the public NIC or LAN network has failed.</p>	<p>Failover. Mirroring continues.</p>	<p>When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.</p>
<p>Public and Private IP or Network Failure</p> <p>May mean that the LAN network, including both private and public NIC connections, has failed.</p>	<p>No interruption of service. No Public LAN access. Mirroring continues.</p>	<p>The site that owned the quorum resource right before the “network partition” remains as owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource.</p>
<p>Lose Network Connection (SAN & LAN), failing both heartbeat and connection to storage</p> <p>May mean that all network and SAN connections are severed, for example if a single pipe is used between buildings for the Ethernet and storage.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The node/site that owned the quorum resource right before the “network partition” remains as owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource. By default MSCS clussvc service will try to auto-start every minute, so after LAN/SAN communication has been re-established, MSCS clussvc will auto-start and will be able to re-join the existing cluster.</p>

Table 13-6 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
Storage Array failure on Site A, or on Site B May mean that a power cord was unplugged, or a storage array failure caused the array to stop responding.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	The campus cluster is divided equally between two sites with one array at each site. Completely failing one storage array should not effect on the cluster or any cluster resources that are currently online. However, you will not be able to move any cluster resources between nodes after this storage failure, because neither node will be able to obtain a majority of disks within the cluster disk group.
Site A failure (power) Means that all access to site A, including server and storage, is lost.	Manual failover.	If the failed site contains the cluster node that owned the quorum resource, then the overall cluster would be offline and cannot be online on the remaining live site without manual intervention.
Site B failure (power) Means that all access to site B, including server and storage, is lost.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	If the failed site did not contain the cluster node that owned the quorum resource, then the cluster would still be alive with whatever cluster resources that were online on that node right before the site failure.

Dealing with a failover situation

In summary, the site scenarios that can occur when there is a cluster server failure include the following:

- If the site not owning the quorum volume and the cluster goes offline, the quorum and data volumes stay online at the other site and other cluster resources stay online or move to that site. Storage Foundation for Windows lets the owning cluster node remain online with 50% ownership of the disks in the quorum group.
- If the site owning the quorum volume goes offline, the remaining site cannot gain control of the quorum volume because it cannot reserve a majority of disks in the quorum group. This is a safeguard to prevent multiple nodes from onlining members of a cluster disk group to which they have access.

Caution: Manual failover of a cluster between two sites should be performed only after coordination between the two sites to ensure that the primary server has in fact failed. If the primary server is still active and you manually import a cluster disk group containing the MSCS quorum to the secondary (failover) server, a split-brain situation occurs. There may be data loss if the split-brain situation occurs because each plex of the mirrored volume may be updated independently when the same disk group is imported on both nodes.

MSCS quorum and quorum arbitration

This section explains the quorum and quorum arbitration in MSCS.

Quorum

The quorum resource maintains the cluster database, as well as critical recovery information, in a recovery log. The quorum resource must be available to all nodes through a SCSI or Fibre Channel bus. With MSCS alone, the quorum disk must be located on a single physical disk. However, with SFW, the quorum disk can be a mirrored volume that spans multiple disks and cluster nodes.

The quorum resource also determines ownership of the cluster. When a node that is controlling the cluster goes offline, other nodes use a challenge/defense protocol to determine which node can have control of the quorum resource and the cluster.

Cluster ownership of the quorum resource

The MSCS challenge/defense protocol uses a low-level bus reset of the SCSI buses between the machines to attempt to gain control of the quorum resource. After a SCSI bus reset, the reservation that each server had been holding on the quorum disk is lost. Each server has about 10 seconds to re-establish that reservation, which would in turn let the other servers know that it is still functioning, even though the other servers would not necessarily be able to communicate with it.

If the active cluster server does not re-establish the SCSI reservation on the quorum resource within the time limit, the applications that were on the server transfer to the server that establishes the SCSI reservation first. The new server servicing the application may now be a bit slower, but clients still get their applications serviced. The IP (Internet Protocol) address and network names move, applications are reconstituted according to the defined dependencies, and clients are still serviced, without any question as to the state of the cluster.

The challenge/defense protocol is more complex when the quorum device is a volume in a Storage Foundation for Windows disk group. For a server to take

ownership of the disk group containing the cluster quorum device, SFW on that server must successfully import the disk group, obtaining SCSI reservations on more than half of its disks. Because a campus cluster configuration has an even number of disks on each site, failover cannot occur automatically. After a site failure, you must use the manual CLI command `vxclus enable` to bring the cluster disk groups online on the secondary node.

The vxclus utility

Storage Foundation for Windows provides the `vxclus` command line utility to allow forcing a failover to the secondary site. The command `vxclus enable` creates an entry in the Registry that enables the cluster disk group to be brought online on a node with a minority of the disks. After you run `vxclus enable`, you can bring the disk group resource online in MSCS Cluster Administrator. After the cluster disk group is brought online, the `vxclus` functionality is disabled.

To bring a cluster online on a node with a minority of the cluster disks

- 1 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.

Caution: When you bring 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 can be corrupted.

- 2 If the cluster service has stopped because of a dynamic quorum resource failure, start the cluster service (`clusvc`).
- 3 Use MSCS Cluster Administrator to bring the cluster disk groups online.

For more information on the `vxclus` utility, see the “Command Line Interface” chapter of the *Storage Foundation Administrator’s Guide*. The `vxclus` utility also provides support for booting from a SAN, but you must have a hardware storage array that supports the capability.

Configuring the network and storage

Use the following procedures to configure the storage hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent MSCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 5 In the Public Status dialog box, on the General tab, click **Properties**.
- 6 In the Public Properties dialog box, on the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.

- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Establishing an MSCS cluster

Before you install SFW, you must install the operating system along with MSCS and then establish an MSCS cluster. After setting up the cluster under MSCS, then you can install SFW and add SFW support with SFW disk groups and volumes.

Note: The steps outlined in this section are general and do not contain specific details. Refer to Microsoft documentation for more complete information.

The tasks for installing the cluster are:

- [“Installing and configuring the operating system and MSCS on Server A”](#) on page 282
- [“Configuring the shared storage and creating a partition for the Cluster quorum disk”](#) on page 283
- [“Creating the first node of the cluster on Server A”](#) on page 283
- [“Installing and configuring the operating system and MSCS on Server B”](#) on page 283
- [“Connecting the two nodes”](#) on page 283
- [“Creating the second node of the cluster on Server B”](#) on page 284
- [“Verifying the cluster configuration”](#) on page 284

Further descriptions of these tasks follow.

Installing and configuring the operating system and MSCS on Server A

This topic summarizes the steps for installing the operating system and configuring the network settings for Server A.

To install and configure the operating system and MSCS on Server A

- 1 Install the Windows Server 2003 operating system on Server A. MSCS is installed as part of the operating system.
- 2 Use the Internet Protocol (TCP/IP) window to identify the static Server A network addresses for the public and private networks in the cluster.
- 3 Make sure a domain is set up that can be used by the cluster nodes, which must be members of the same domain.

- 4 Select **Administrative Tools > Active Directory > Users and Computers** and set up a cluster account for the cluster. Microsoft recommends having a separate user account under which the cluster can run.

Configuring the shared storage and creating a partition for the Cluster quorum disk

Configuring the shared storage and creating a partition for the cluster quorum disk, consists of the following tasks:

- Configure the disks for the storage array attached to Server A.
- Use **Disk Management** to create a partition for the cluster quorum disk on a basic disk that will be used as the quorum disk when the first node of the cluster is created.
Microsoft recommends 500 MB as the partition size and includes the entire disk as a cluster resource.

Creating the first node of the cluster on Server A

Create the first node of the cluster on Server A. Refer to the Microsoft documentation for details.

After you establish the cluster on Server A, make sure that you can see the storage array's disks from Server A.

Installing and configuring the operating system and MSCS on Server B

Repeat the same installation steps for Server B as you used for Server A.

See [“Installing and configuring the operating system and MSCS on Server A”](#) on page 282.

Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so now MSCS controls the cluster storage on Server A, and the operating system cannot access both nodes of the storage at the same time.

To connect the two nodes

- 1 Connect the corresponding cables between the three network cards on the two sites.
- 2 Connect the two switches at the two sites through the storage interconnect.
- 3 Test the connectivity between the two sites, by doing the following:

- Test the IP addresses of all the network adapter cards in the cluster.
- Bring up the command window and type `ping ipaddress`, where the *ipaddress* is the corresponding network adapter in the other node.

Creating the second node of the cluster on Server B

Create the second node of the cluster on Server B. Refer to the Microsoft documentation for details.

Verifying the cluster configuration

After the configuration is complete, use the following procedure to verify failover.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.

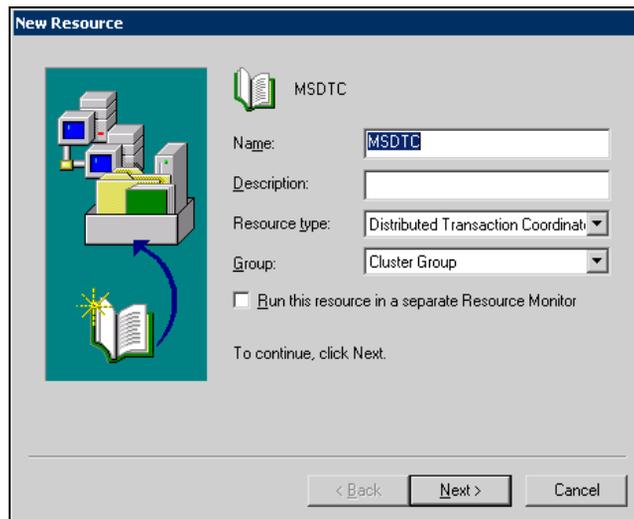
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Creating the MSDTC resource (Exchange 2003 only)

Before you install Exchange 2003, create the MSDTC resource. You can create this resource now or just before installing Exchange.

To create the MSDTC resource

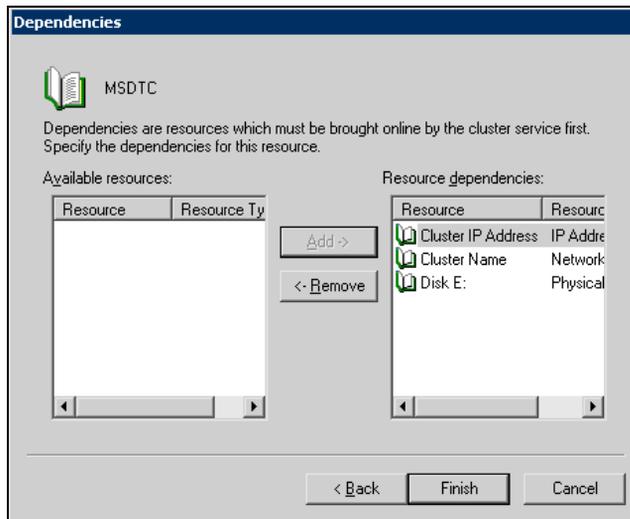
- 1 In Cluster Administrator (**Start > Administrative Tools > Cluster Administrator**), right-click **Cluster Group**, click **New**, and click **Resource**.
- 2 In the New Resource dialog box, specify the following options and then click **Next**.



Name	Type a name for the MSDTC resource.
Description	If necessary, type a description of the resource
Resource type	Click Distributed Transaction Coordinator .

- 3 In the Possible Owners dialog box, all the nodes in the cluster are listed as possible owners. Click **Next**.

- 4 In the Dependencies dialog box, select the cluster IP address, cluster name, and physical disk resources from the Available Resources list, add them to the Resource dependencies list, and click **Finish**.



- 5 Click **OK**.
- 6 In the left pane, expand the Groups icon.
- 7 Click **Cluster Group**.
- 8 Right-click **Bring Online**.
The state changes to online.

Installing SFW

This section assumes you are running an MSCS cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the MSCS cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
 See “[Pre-installation tasks](#)” on page 287.
- Installing the product
 See “[Installing Veritas Storage Foundation for Windows](#)” on page 289.
- Performing post-installation tasks
 See “[Post-installation tasks](#)” on page 292.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
 See “[Changing the driver signing options](#)” on page 287.
- Moving the Online Groups
 See “[Moving the online groups](#)” on page 288.

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The table below describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 13-7 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed

Table 13-7 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
 If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
 If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a MSCS configuration.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.

The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements. The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).

The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.

- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.
 - Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.
If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.
To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.
Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/ Failover Cluster** option.
 - Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
 - To add a computer for installation, provide the name of the computer in the System Name box.
OR
If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.
 - To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
 - To know the verification status and other information of the added computer, click the information icon.
 - To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

- **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

- **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.

It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.

Click **Install** to install the software.

- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups
See “[Moving the online groups](#)” on page 292.
- Completing the SFW Installation
See “[Completing the SFW installation](#)” on page 292.
- Resetting the driver signing options
See “[Resetting the driver signing options](#)” on page 293.

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.
- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 286.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Creating disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes on the cluster storage, which in a campus cluster consists of at least two storage arrays.

Before you create disk groups and volumes, consider the following items:

- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs which depend on the traffic load
- The disk groups and number of disks on each site

Note: For campus clusters, each disk group must contain an equal number of disks on each site.

- Types of volumes required and location of the plex of each volume in the storage array

Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Create two or more dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum.

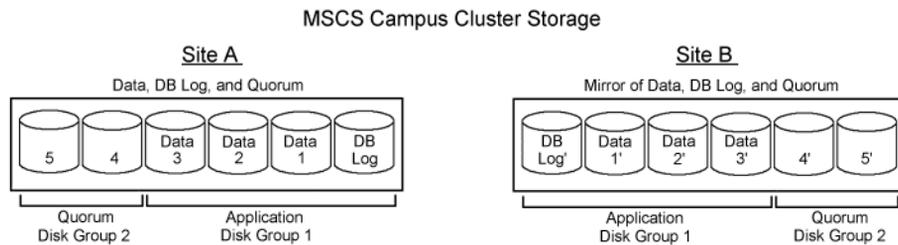
For the Microsoft Exchange application data files, you could create a separate disk group for each storage group. It is best to separate data files from log files and place them in separate volumes. For example, you might create an Exchange disk group for the first storage group of the first Exchange Virtual Server, EVS1_SG1_DG, with the following two volumes:

- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

Figure 13-3 shows a typical MSCS campus cluster setup of disks. This example has only one application disk group that spans the storage arrays at both sites. The data and database log on Site A are mirrored to Site B. Each mirrored volume can have more than two disks, but must have an even number, such as four. All the application data could be in one large mirrored volume with multiple disks, but the same number of disks are required on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

In [Figure 13-3](#), a four-way mirror for the quorum volume provides additional redundancy. The minimum configuration is a two-way mirror. If possible, use small disks for the quorum volume. Microsoft recommends 500 MB for the quorum volume.

Figure 13-3 MSCS campus cluster disks and disk groups example



Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site, and that each volume is a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 296
- [“Creating a volume”](#) on page 298

Considerations when creating new volumes

- For campus clusters, when you create a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored for the new volumes. Striped mirrored gives you better performance compared to concatenated.

When selecting striped mirrored, select two columns to stripe one enclosure that is mirrored to the second enclosure.

- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and, if prompted, select a profile.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. If prompted, provide the user name, password, and domain.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.

The internal names for the disks which the current system can access for available storage are displayed, with names **Harddisk1**, **Harddisk2**, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

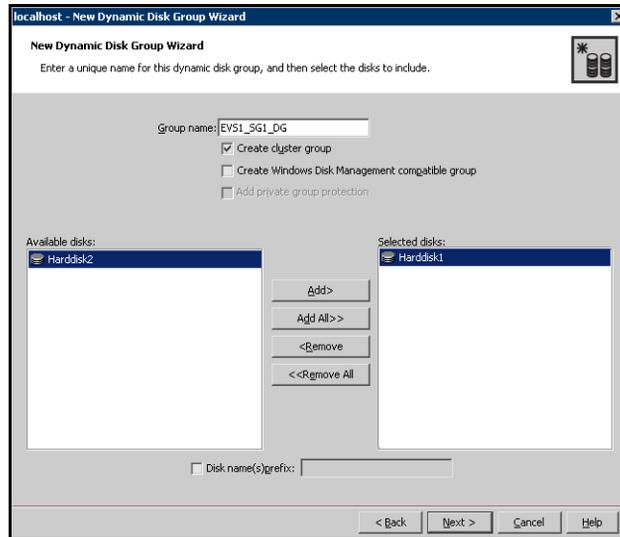
Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier.
 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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.

- 8 Click **Finish** to create the new disk group.

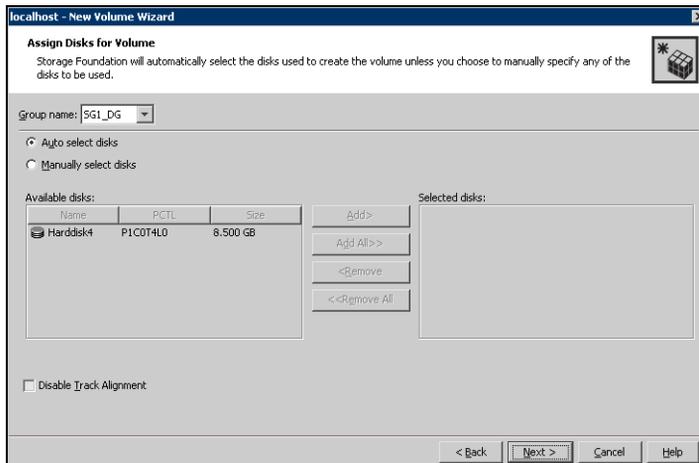
Proceed to create the appropriate volumes on each disk.

Creating a volume

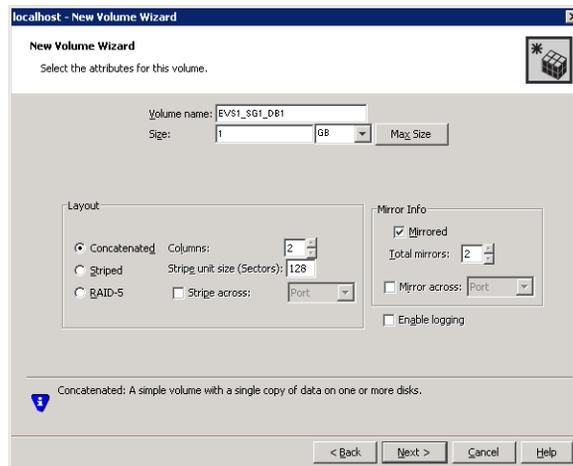
Use the following procedure to create dynamic volumes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.

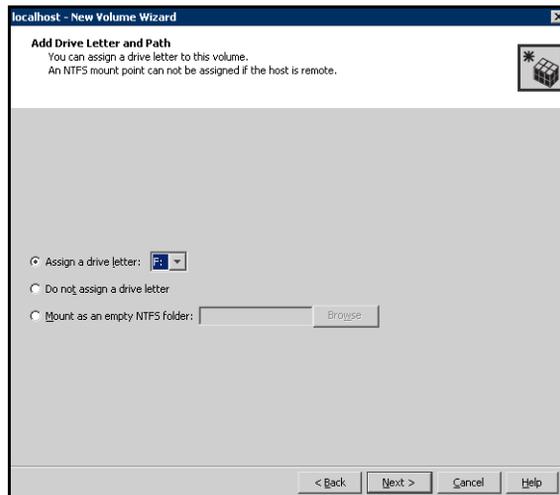


- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3COT2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.
- 9 Specify the volume attributes.



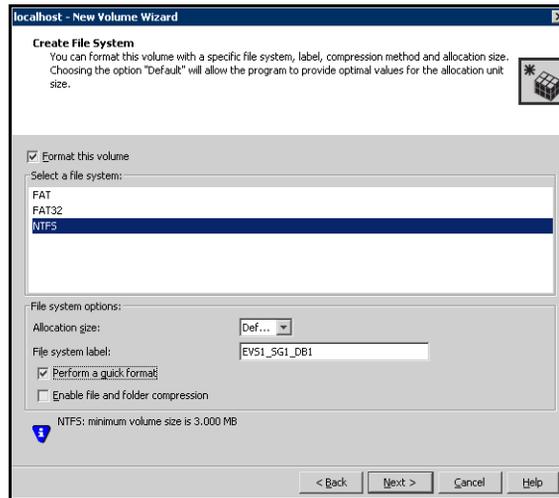
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.

- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.

13 Click **Finish** to create the new volume.

14 Repeat these steps to create the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) For Exchange 2003, you could also create an MTA volume (EVS1_SG1_MTA).

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Implementing a dynamic quorum resource

One of the key advantages of using SFW with MSCS is that you can create a mirrored quorum resource that adds fault tolerance to the quorum, thus protecting the cluster from failure if the disk that the quorum is on fails. In the following procedure, you transfer the cluster's quorum resource from a physical disk resource to a mirrored dynamic quorum resource. The tasks for creating a mirrored quorum resource are:

- [“Creating a dynamic cluster disk group for the quorum, mirrored”](#) on page 303
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 304
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 305

Creating a dynamic cluster disk group for the quorum, mirrored

If you have not already done so, use SFW to create a dynamic disk group for the quorum. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using four small disks for the mirrored quorum for additional redundancy.

If possible, use small disks, because the disk group will only be used for the quorum volume, which Microsoft recommends to be 500 MB. To create a four-way mirrored volume in the New Volume wizard, select the **Concatenated** layout, click the **Mirrored** checkbox, and specify four mirrors. For full details on creating cluster disk groups and volumes, see:

[“Creating disk groups and volumes”](#) on page 294.

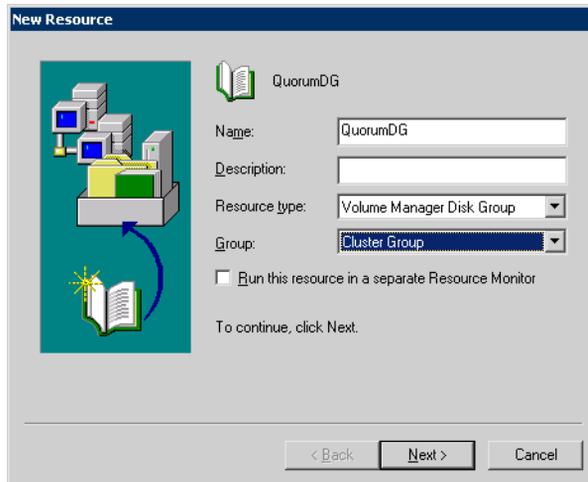
Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

Making the quorum cluster disk group an MSCS resource

The dynamic cluster disk group that you prepared for the quorum needs to be added as a resource to the default Cluster Group in MSCS. Complete this step now if you have not done it earlier.

To make the quorum disk group an MSCS resource

- 1 Verify that the Cluster Group is online on the same node where you created the cluster disk group for the quorum.
- 2 Right-click the Cluster Group and select **New > Resource**. The New Resource window appears.

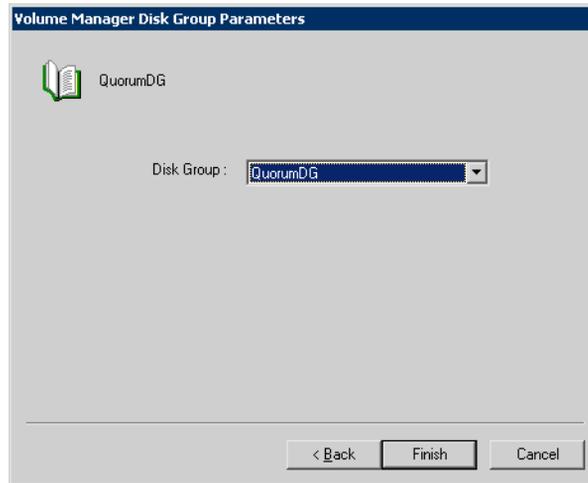


- 3 On the New Resource, window, do the following:
 - Specify a name for the disk group resource in the **Name** field, such as “QuorumDG.”
 - If necessary, you can add a description of the resource in the **Description** field.
 - Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.

Note: The resource name has not been changed to Storage Foundation Disk Group.

- Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
- Click **Next**.

- 4 On the Possible Owners screen, by default all the nodes in the cluster are listed as possible owners. Click **Next**.
- 5 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)
- 6 Make sure the appropriate SFW quorum cluster dynamic disk group is selected from the drop-down list for the resource, and click **Finish** to complete the operation.



Changing the quorum resource to the dynamic mirrored quorum resource

Use Cluster Administrator to change the quorum resource from a physical disk resource to a dynamic disk quorum resource.

To change the quorum resource to the dynamic mirrored quorum resource

- 1 From Cluster Administrator, right-click the cluster name in the tree view to bring up its context menu.
- 2 Select **Properties**, which displays the Properties window.
- 3 Click the **Quorum** tab of the Properties window.
- 4 Select the name of the dynamic quorum disk group as the resource to be used for the quorum resource.
- 5 Click **OK**.

Adding a Volume Manager Disk Group resource for Exchange 2007 installation

Before Exchange 2007 installation you add a Volume Manager Disk Group resource to the Cluster Group so that you can install the First Storage Group database files to an SFW dynamic volume.

After Exchange installation you will move that resource from the Cluster Group to the Exchange group that is created during Exchange 2007 installation.

To add a Volume Manager Disk Group resource for Exchange 2007

- 1 Verify that the Cluster Group is online on the node.
- 2 Right-click on the Cluster Group and select **New > Resource**. The New Resource window appears.
- 3 On the New Resources window, do the following:
 - Specify a name for the disk group resource in the **Name** field.
 - If required, add a description about the resource in the **Description** field.
 - Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.
 - If necessary, use the drop-down list to select the appropriate MSCS group; the group should already be selected.
 - Generally, make sure that **Run this resource in a separate Resource Monitor** is not checked.
 - Click **Next**.
- 4 On the Possible Owners screen, all the nodes in the cluster are listed as possible owners by default. Click **Next**.
- 5 On the Dependencies screen, click **Next**. (You do not need to set any dependencies for a disk group resource.)
- 6 Make sure the appropriate SFW cluster disk group is selected from the drop-down list for the resource, and click **Finish**.

Installing the application on the cluster nodes

You must install the application program files on the same local drive of all the cluster nodes.

For any specific requirements for the application in an MSCS environment, see the Microsoft documentation.

Checklist for installing the application on the first node

- Applications may have built-in procedures for running on a cluster. Consult the application documentation to determine whether these procedures are available.
- Some applications, such as Microsoft Exchange Server and Microsoft SQL Server, install on both nodes at once.
- All nodes of the clustered application must share the same virtual name and IP address.
- When you install Exchange 2007, remember not to accept the default location for the database. Instead, click to browse to the dynamic volume that was prepared previously.

Checklist for installing the application on the second node

- To install the application on the second node, move the cluster resources to the second node.
- Make sure that the shared volumes, when accessed on the second node, have the corresponding drive letters or mount points that they had when accessed from the first node.
- If you are installing a database, you may need to stop the database service on the first node while the shared disks are being manipulated by the installation on the second node. Then restart the service after the application is installed.

To add or change a drive letter or mount point

- 1 In VEA, right-click on the volume for which the drive letter will be added or changed.
- 2 Select **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths window, add or change a drive letter, or add or change a mount point.
 - To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter.

- To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Change the drive letter.
- To add a mount point, click the **Add** radio button, click the **Mount as an empty NTFS folder** radio button, browse to select an empty folder or click the **New Folder** button to create a new folder.

Note: A mount point is also referred to as a “drive path.”

- To change a mount point, you must remove it and add it again. (See the bullet above). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.
- Click **OK**.

Setting up a group for Exchange 2003 in MSCS

Use MSCS Cluster Administrator to set up an application group for Exchange 2003. Add to the Exchange group the IP Address resource, the Network Name resource, and the Microsoft Exchange System Attendant resource.

Before you create the application resource, make sure that all the other resources that you created are online, including the disk group resource and any additional application resources.

For help on creating the application resource and additional resources that may be required, see the application documentation.

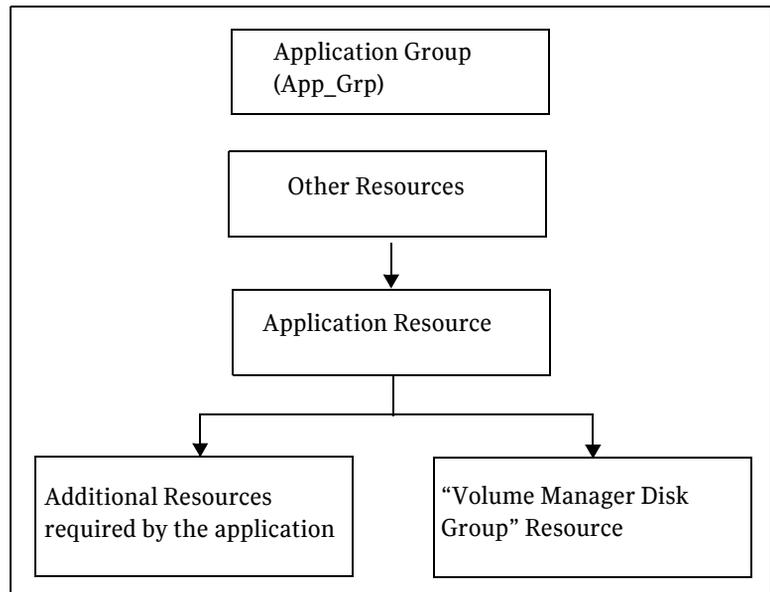
You cannot add the application resource until after the application is installed on both nodes.

Add Volume Manager Disk Group resources for any SFW disk groups you created for Exchange.

When you add resources, make sure that you configure the appropriate resource dependencies. You must configure the application resource (in the case of Exchange 2003, the System Attendant resource) to depend on the Volume Manager Disk Group resource(s).

Figure 13-4 shows the application group dependencies.

Figure 13-4 Application group dependencies



Adding Volume Manager Disk Group resources to the Exchange 2007 group

For Exchange 2007 you add a Volume Manager Disk Group resource to the Cluster Group before installing Exchange. By adding the resource before installation, you can install the First Storage Group to a dynamic volume.

See “[Adding a Volume Manager Disk Group resource for Exchange 2007 installation](#)” on page 306.

Installing Exchange 2007 in the MSCS cluster automatically sets up a group for Exchange 2007 with the required resources for Exchange 2007.

Use the Cluster Administrator to move the Volume Manager Disk Group resource that you added for Exchange 2007 from the Cluster Group to the Exchange group. After doing so, set the Exchange database resource for the First Storage Group to depend on the Volume Manager Disk Group resource.

If you have created additional cluster disk groups for Exchange storage groups, add a Volume Manager Disk Group resource for each cluster disk group to the Exchange group. Set the appropriate dependencies.

Moving Exchange databases and logs to dynamic volumes

If you plan to use the First Storage Group in your Exchange cluster configuration, ensure that the First Storage Group database and log are located on the dynamic volumes.

Choose the appropriate procedure to move Exchange databases and logs, depending on whether you are configuring Exchange 2003 or 2007.

Moving Exchange 2003 databases and logs

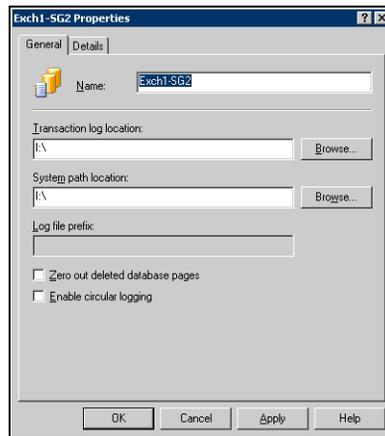
For Exchange 2003, set the path for the transaction log location and system path location fields to point to the log volume. Set the path for the Exchange database and Exchange streaming database files to point to the appropriate database volumes. You use the Exchange System Manager to specify the locations.

Optional steps for creating a new storage group and mailbox stores are included in the procedure.

To point the Exchange 2003 databases and log paths to the SFW volumes

- 1 Click **Start > All Programs > Microsoft Exchange > System Manager** to open the Exchange 2003 System Manager.

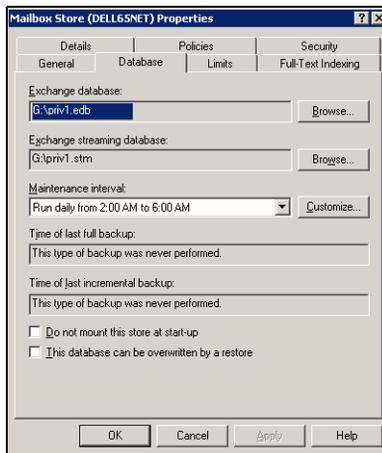
- 2 In the appropriate Administrative Group, expand **Servers** and expand the appropriate Exchange server.
- 3 Choose one of the following:
 - Right-click an existing storage group, for example SG2, and click **Properties**.
 - To create a new storage group, right-click the appropriate Exchange server, click **New** and click **Storage Group**.
- 4 Complete the following on the Properties sheet and click **OK**:



Name	For a new storage group, enter the name storage group name (for example, SG2).
Transaction log location	Click the Browse button and set the transaction log path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \.
System path location	Click the Browse button and set the path to the drive letter or mount point of the volume created for the transaction logs of the storage group, for example I : \.
	The paths for the Transaction log location and System path location should be the same.
Enable circular logging	Make sure that the Enable circular logging check box is not checked.

- 5 Right-click on the storage group and choose one of the following:

- For an existing storage group, right-click on a mailbox store and click **Properties**.
 - For a new storage group, click **New** and click **Mailbox Store**.
- 6 Choose one of the following:
- For an existing storage group, proceed to [step 7](#).
 - For a new storage group, in the **General** tab of the Properties sheet enter the name of the new mailbox store (for example, SG2-DB1).
- 7 Click the **Database** tab, set the paths for the .edb and .stm files for the database as follows, and click **OK**:



Exchange database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing EDB files (for example, G: \).

Exchange streaming database Click the **Browse** button and set the path to the drive letter or mount point of the volume created for storing STM files (for example, G: \).

- 8 Click **Yes** to mount the store.
- 9 Repeat [step 5](#) through [step 8](#) to create or set the paths for other mailbox stores. For example, create another mailbox store mounted on the H: drive, SG2-DB2.

Moving Exchange 2007 databases and logs

During Exchange 2007 installation, the First Storage Group database and log are installed on the same volume. You move the log to a separate dynamic volume after installation. If you need to move any other Exchange 2007 databases or logs, use the same procedure.

For each Exchange 2007 storage group, set the log files path and system files path to point to the log volume. For each database, set the database path to point to the appropriate database volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new storage groups or databases or change the locations for existing storage groups or databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

The following procedures can be used to change paths for existing storage groups and databases.

See the Microsoft Exchange 2007 documentation for additional information on creating new storage groups and databases.

To use the Exchange Management Console to change log file and system file locations for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the storage group for which you want to change the log file location.
- 4 In the work pane, right-click the storage group for which you want to change the log file location and click **Move Storage Group Path**. The Move Storage Group Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the log files, and create a folder for the log files. Repeat for the system files. The volume should be the same for both.
- 6 Click **Move**. A warning appears that all databases in the storage group must be temporarily dismounted, which will make them inaccessible to any user. To continue, click **Yes**.

- 7 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Storage Group Path wizard.

To use the Exchange Management Console to change the database file location for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the database for which you want to change the database file location.
- 4 In the work pane, right-click the desired database and click **Move Database Path**. The Move Database Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the database file, and create a folder for it as needed.
- 6 Click **Move**. A warning appears indicating that the database must be temporarily dismounted, which will make it inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the database files were moved successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Verifying the cluster configuration

After you complete the configuration, verify that failover occurs as desired.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use the Cluster Administrator to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator**.
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node use **Move Group**.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.
- 2 Open the Cluster Administrator. Click **Start > All Programs > Administrative Tools > Cluster Administrator** from any node in the cluster.
- 3 In the Cluster Administrator, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move all the resource groups back to the original node, restart the node you shut down in [step 1](#) and use **Move Group** to move all the resource groups.

Deploying SFW with Microsoft failover clustering and Exchange in a campus cluster

This chapter covers the following topics:

- [Tasks for a new Exchange, SFW and Microsoft failover clustering installation in a campus cluster \(Windows Server 2008\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the network and storage](#)
- [Establishing a Microsoft failover cluster](#)
- [Installing SFW](#)
- [Creating disk groups and volumes](#)
- [Implementing a dynamic quorum resource](#)
- [Adding a Volume Manager Disk Group resource for Exchange 2007 installation](#)
- [Installing the application on the cluster nodes](#)
- [Adding the Volume Manager Disk Group resources to the Exchange group](#)
- [Setting the database dependency on the disk group resource](#)
- [Moving the Exchange databases and logs to the dynamic volumes](#)
- [Verifying the cluster configuration](#)

Tasks for a new Exchange, SFW and Microsoft failover clustering installation in a campus cluster (Windows Server 2008)

On Windows Server 2008, you can install and configure Veritas Storage Foundation for Windows with Microsoft failover clustering and Microsoft Exchange Server 2007. This chapter presents a campus clustering example using a two-node cluster.

This chapter presents a Microsoft failover clustering example with a two-node campus cluster. This chapter describes the deployment on Windows Server 2008.

The table below outlines the high-level objectives and the tasks for each objective:

Table 14-1 Task list for deploying SFW with Exchange 2007 and Microsoft failover clustering in a campus cluster

Objectives	Tasks
“ Reviewing the requirements ” on page 320	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites.
“ Reviewing the configuration ” on page 323	<ul style="list-style-type: none"> ■ Review the configuration requirements. ■ Overview of a campus cluster using Microsoft clustering and recovery scenarios.
“ Configuring the network and storage ” on page 331	<ul style="list-style-type: none"> ■ Install and configure the hardware for each node in the cluster. ■ Verify the DNS settings and binding order for all systems.
“ Establishing a Microsoft failover cluster ” on page 334	<ul style="list-style-type: none"> ■ Enable the Microsoft failover clustering feature. ■ Ensure that you have met the hardware requirements for a failover cluster. ■ Run the Microsoft wizard to validate the configuration. ■ Use Failover Cluster Management to create the first node of the cluster. ■ Create the second node of the cluster. ■ Connect the two nodes.
“ Installing SFW ” on page 336	<ul style="list-style-type: none"> ■ Install SFW on Node A (Node B active). ■ Install SFW on Node B (Node A active).

Table 14-1 Task list for deploying SFW with Exchange 2007 and Microsoft failover clustering in a campus cluster (Continued)

Objectives	Tasks
<p>“Creating disk groups and volumes” on page 342</p>	<ul style="list-style-type: none"> ■ In SFW on Node A, create two or more dynamic cluster disk groups on the storage, one or more for the application data files and one for the mirrored quorum.
<p>“Implementing a dynamic quorum resource” on page 351</p>	<ul style="list-style-type: none"> ■ If not done earlier, in the VEA, create a dynamic disk group for the quorum with a mirrored volume. ■ In Failover Cluster Management, create a quorum group and add to it the Volume Manager Disk Group resource for the quorum disk group. ■ Change the quorum resource to the dynamic mirrored quorum resource.
<p>“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 354</p>	<ul style="list-style-type: none"> ■ Add a Volume Manager Disk Group resource for the SFW disk group that was created for the First Storage Group. You add it to the quorum group and then later move it to the Exchange group. This allows the First Storage Group to be installed to a dynamic disk.
<p>“Installing the application on the cluster nodes” on page 355</p>	<ul style="list-style-type: none"> ■ Follow Microsoft instructions for installing the Exchange clustered mailbox role. Install files relating to the data and logs on the shared storage.
<p>“Adding the Volume Manager Disk Group resources to the Exchange group” on page 356</p>	<ul style="list-style-type: none"> ■ Moving the Volume Manager Disk Group resource for the First Storage Group to the Exchange 2007 group that was created during Exchange installation. Adding any additional SFW disk groups to the Exchange group as additional Volume Manager Disk Group resources.
<p>“Setting the database dependency on the disk group resource” on page 356</p>	<ul style="list-style-type: none"> ■ For each database, set the dependency on the Volume Manager Disk Group resource that was created for it.
<p>“Moving the Exchange databases and logs to the dynamic volumes” on page 357</p>	<ul style="list-style-type: none"> ■ The First Storage Group database and log are installed to one volume; move the log to a separate dynamic volume.

Table 14-1 Task list for deploying SFW with Exchange 2007 and Microsoft failover clustering in a campus cluster (Continued)

Objectives	Tasks
“ Verifying the cluster configuration ” on page 359	<ul style="list-style-type: none">■ Verify the cluster configuration by either moving all the resource groups from one node to another or by simulating a failover by shutting down the active cluster node.

Reviewing the requirements

Reviewing the requirements and the configuration allows you to gain an overall understanding of the configuration and its requirements.

See the following topics:

- [Supported software](#)
- [System requirements](#)
- [Disk space requirements](#)

Supported software

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW and Microsoft clustering on Windows Server 2008:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2008 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 14-2 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2008 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

System requirements

- One CD-ROM drive accessible to each system on which you are installing Microsoft clustering.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access the storage.
- Microsoft clustering requires at least two network adapters per system (one network adapter to connect each system to the public network and one network adapter for the private network on each system). Symantec recommends having two adapters for the private network and routing each

private network adapter through a separate hub or switch to avoid single points of failure.

- Refer to application documentation to determine disk space requirements for your application.
- Each system requires 1 GB of RAM for SFW.
- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- The configuration requires two sites with a storage array for each site, with an equal number of disks at each site for the mirrored volumes.
- Interconnects between the clusters are required for the storage and the network.
- Systems to be clustered must be configured as part of a Windows Server 2008 domain. Each system in a cluster with Microsoft failover clustering must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and Microsoft clustering software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. Six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.

Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 14-3 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB

Table 14-3 Disk space requirements

Installation options	Install directory/drive
SFW + all options	950 MB
Client components	354 MB

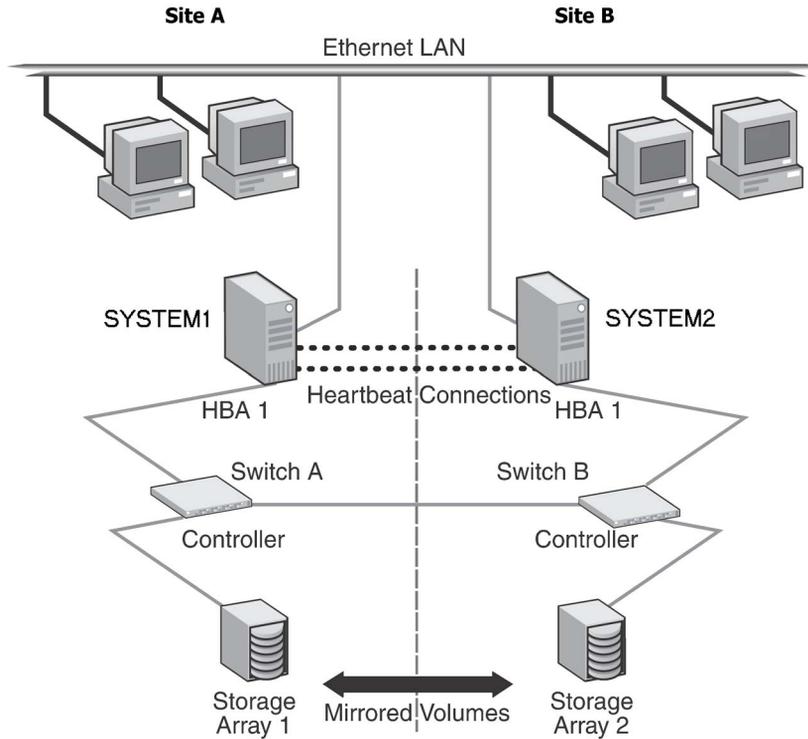
Note: Plan for an equal number of disks on the two sites, because each disk group should contain the same number of disks on each site.

Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

[Figure 14-1](#) illustrates an example campus cluster configuration.

Figure 14-1 Campus clustering with Microsoft clustering configuration example



The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array and contains mirrored data of the storage on the other array. Each disk group should contain the same number of disks on each site for the mirrored volumes.

Microsoft clustering uses the quorum architecture, where the cluster database resides in the quorum resource. If you are using Microsoft clustering, adding SFW to the configuration protects the quorum disk from being a single point of failure in the cluster because SFW provides dynamic volumes and software mirroring of the quorum device. To avoid a single point of failure, set up the quorum as a dynamic mirrored device. This example includes the dynamic mirrored quorum and requires setting up two or more dynamic cluster disk groups in SFW— one or more cluster disk groups for the application and data and one for the dynamic mirrored quorum.

The example configuration does not include DMP. For instructions on how to add DMP to a clustering configuration, see *Veritas Storage Foundation and High Availability Solutions, Solutions Guide*.

When you are installing SFW and Microsoft clustering together, remember the following:

- A cluster using Microsoft clustering must be running to install SFW. You need to set up the hardware and install the operating system and Microsoft clustering on all systems and establish the failover cluster before installing SFW.

Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Use a “rolling install” procedure to install SFW first on the inactive cluster node. Then move the cluster resources to the other node and install on the now inactive node.

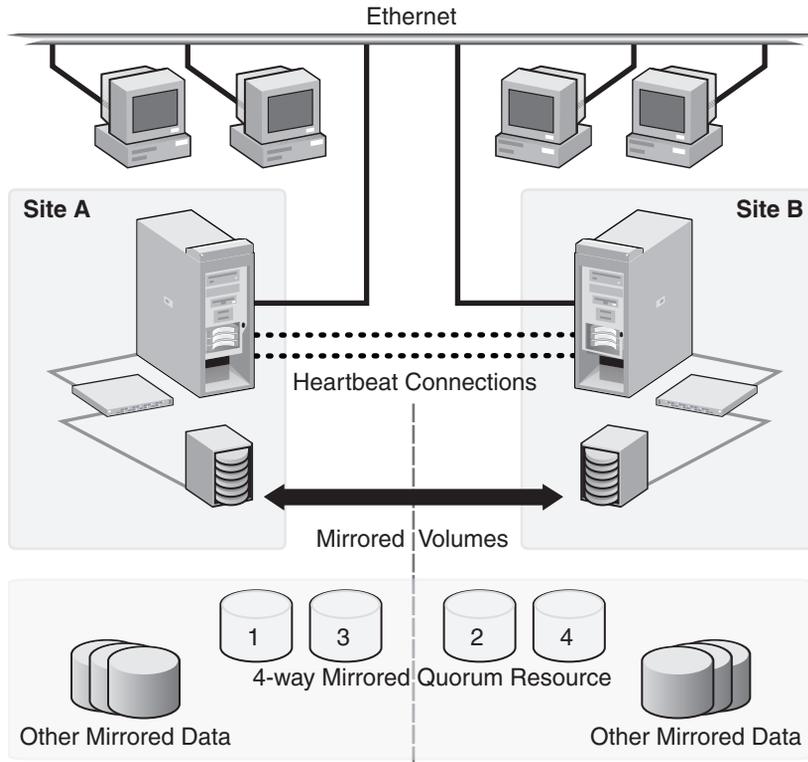
Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- After SFW is installed, create one or more cluster disk groups with SFW and set up the volumes for your application. At the same time, you can create the mirrored volume for the dynamic quorum resource.
- SFW allows you to add fault tolerance to your data volumes. Mirroring of log volumes is recommended, and a mirrored striped RAID layout is recommended for your data volumes. SFW also offers multiple disk groups, multiple mirrors, capacity management and Automatic Volume Growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, Dynamic Multi-pathing, and enhanced snapshot capabilities with FlashSnap.

Overview of campus clustering with Microsoft clustering

Figure 14-2 on page 326 shows a campus cluster configuration with Microsoft clustering. It features mirrored storage across clusters and a mirrored quorum resource. The figure shows a four-way mirrored quorum that has an extra set of mirrors for added redundancy. Although a campus cluster setup with Microsoft clustering can work without Storage Foundation for Windows, SFW provides key advantages over using Microsoft clustering alone. Through a dynamic mirrored volume that functions on multiple disks across multiple sites, SFW protects the quorum resource in the cluster from being the single point of failure in the cluster.

Figure 14-2 Typical campus clustering configuration with Microsoft clustering



Most customers use hardware RAID to protect the quorum disk, but that will not work when a natural disaster takes down the primary node and its attached storage. If the quorum resource is lost to the cluster, the cluster fails, because none of the cluster servers can gain control of the quorum resource and ultimately the cluster. Microsoft clustering alone cannot provide fault tolerance to the quorum disk.

Campus cluster failure with Microsoft clustering scenarios

This section focuses on the failure and recovery scenarios with a campus cluster with Microsoft clustering and SFW installed.

For information about the quorum resource and arbitration in Microsoft clustering, see

[“Microsoft clustering quorum and quorum arbitration”](#) on page 330.

Table 14-4 lists failure situations and the outcomes that occur:

Table 14-4 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p>Application fault</p> <p>May mean the services stopped for an application, a NIC failed, or a database table went offline.</p>	Failover	If the services stop for an application failure, the application automatically fails over to the other site.
<p>Server failure (Site A)</p> <p>May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.</p>	Failover	Assuming a two-node cluster pair, failing a single node results in a cluster failover. Service is temporarily interrupted for cluster resources that are moved from the failed node to the remaining live node.
<p>Server failure (Site B)</p> <p>May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.</p>	No interruption of service.	Failure of the passive site (Site B) does not interrupt service to the active site (Site A).
<p>Partial SAN network failure</p> <p>May mean that SAN fiber channel cables were disconnected to Site A or Site B Storage.</p>	No interruption of service.	Assuming that each of the cluster nodes has some type of Dynamic Multi-pathing (DMP) solution, removing one SAN fiber cable from a single cluster node should not effect any cluster resources running on that node, because the underlying DMP solution should seamlessly handle the SAN fiber path failover.
<p>Private IP Heartbeat Network Failure</p> <p>May mean that the private NICs or the connecting network cables failed.</p>	No interruption of service.	With the standard two-NIC configuration for a cluster node, one NIC for the public cluster network and one NIC for the private heartbeat network, disabling the NIC for the private heartbeat network should not effect the cluster software and the cluster resources, because the cluster software simply routes the heartbeat packets through the public network.

Table 14-4 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
<p>Public IP Network Failure May mean that the public NIC or LAN network has failed.</p>	<p>Failover. Mirroring continues.</p>	<p>When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.</p>
<p>Public and Private IP or Network Failure May mean that the LAN network, including both private and public NIC connections, has failed.</p>	<p>No interruption of service. No Public LAN access. Mirroring continues.</p>	<p>The site that owned the quorum resource right before the “network partition” remains the owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource.</p>
<p>Lose Network Connection (SAN & LAN), failing both heartbeat and connection to storage May mean that all network and SAN connections are severed; for example, if a single pipe is used between buildings for the Ethernet and storage.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The node/site that owned the quorum resource right before the “network partition” remains the owner of the quorum resource, and is the only surviving cluster node. The cluster software running on the other cluster node self-terminates because it has lost the cluster arbitration for the quorum resource. By default, the Microsoft clustering clussvc service tries to auto-start every minute, so after LAN/SAN communication has been re-established, the Microsoft clustering clussvc auto-starts and will be able to re-join the existing cluster.</p>
<p>Storage Array failure on Site A, or on Site B May mean that a power cord was unplugged, or a storage array failure caused the array to stop responding.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>The campus cluster is divided equally between two sites with one array at each site. Completely failing one storage array should have no effect on the cluster or any cluster resources that are online. However, you cannot move any cluster resources between nodes after this storage failure, because neither node will be able to obtain a majority of disks within the cluster disk group.</p>

Table 14-4 List of failure situations and possible outcomes

Failure Situation	Outcome	Comments
Site A failure (power) Means that all access to site A, including server and storage, is lost.	Manual failover.	If the failed site contains the cluster node that owned the quorum resource, then the overall cluster is offline and cannot be online on the remaining live site without manual intervention.
Site B failure (power) Means that all access to site B, including server and storage, is lost.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	If the failed site did not contain the cluster node that owned the quorum resource, then the cluster is still alive with whatever cluster resources that were online on that node right before the site failure.

Dealing with a failover situation

In summary, the site scenarios that can occur when there is a cluster server failure include the following possibilities:

- If the site not owning the quorum volume and the cluster goes offline, the quorum and data volumes stay online at the other site and other cluster resources stay online or move to that site. Storage Foundation for Windows allows the owning cluster node to remain online with 50% ownership of the disks in the quorum group.
- If the site owning the quorum volume goes offline, the remaining site cannot gain control of the quorum volume because it cannot reserve a majority of disks in the quorum group. This is a safeguard to prevent multiple nodes from online members of a cluster disk group to which they have access.

Caution: Manual failover of a cluster between two sites should be performed only after coordination between the two sites to ensure that the primary server has actually failed. If you manually import a cluster disk group containing the Microsoft clustering quorum to the secondary (failover) server when the primary server is still active, this causes a split-brain situation. If the split-brain situation occurs, you may lose data because each plex of the mirrored volume may be updated independently when the same disk group is imported on both nodes.

Microsoft clustering quorum and quorum arbitration

This section provides an explanation of the quorum and quorum arbitration in Microsoft clustering.

Quorum

The quorum resource maintains the cluster database, as well as critical recovery information, in a recovery log. The quorum resource has to be available to all nodes through a SCSI or Fibre Channel bus. With Microsoft clustering alone, the quorum disk must be located on a single physical disk. However, with SFW, the quorum disk can be a mirrored volume that spans multiple disks and cluster nodes.

The quorum resource also determines ownership of the cluster. When a node that is controlling the cluster goes offline, other nodes use a challenge/defense protocol to determine which node can have control of the quorum resource and the cluster.

Cluster ownership of the quorum resource

The Microsoft clustering challenge/defense protocol uses a low-level bus reset of the SCSI buses between the machines to attempt to gain control of the quorum resource.

After a SCSI bus reset, the reservation that each server had been holding on the quorum disk is lost. Each server then has roughly 10 seconds to re-establish that reservation, which would in turn let the other servers know that it is still functioning, even though the other servers would not necessarily be able to communicate with it.

If the active cluster server does not re-establish the SCSI reservation on the quorum resource within the time limit, all applications that were on the server will then transfer to the server that establishes the SCSI reservation first. The new server servicing the application may now be a bit slower, but clients will still get their applications serviced. The IP (Internet Protocol) address and network names will move, applications will be reconstituted according to the defined dependencies, and clients will still be serviced, without any question as to the state of the cluster.

The challenge/defense protocol is more complex when the quorum device is a volume in a Storage Foundation for Windows disk group. For a server to take ownership of the disk group containing the cluster quorum device, SFW on that server must successfully import the disk group, obtaining SCSI reservations on more than half of its disks. Because a campus cluster configuration has an even number of disks on each site, failover cannot occur automatically. The manual

CLI command, `vxclus enable` must be used to bring the cluster disk groups online on the secondary node after a site failure.

The vxclus utility

Storage Foundation for Windows provides the `vxclus` command line utility to allow forcing a failover to the secondary site. The command `vxclus enable` creates an entry in the Registry that enables the cluster disk group to be brought online on a node with a minority of the disks. Once `vxclus enable` is executed, you can bring the disk group resource online in Failover Cluster Management. After the cluster disk group is brought online, the `vxclus` functionality is disabled.

To bring a cluster online on a node with a minority of the cluster disks

- 1 Use the following `vxclus` command for each disk group on your cluster node:

```
vxclus enable -g<DynamicDiskGroupName>
```

You will be asked to confirm the use of this command.

Caution: 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.

- 2 If the cluster service has stopped because of a dynamic quorum resource failure, start the cluster service (`clussvc`).
- 3 Then, using Failover Cluster Management, bring the cluster disk groups online.

For more information on the `vxclus` utility, see the “Command Line Interface” chapter of the *Storage Foundation Administrator’s Guide*. The `vxclus` utility also provides support for booting from a SAN, but you must have a hardware storage array that supports the capability.

Configuring the network and storage

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.

- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent the Microsoft cluster from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- 3 Use independent hubs or switches for the private heartbeats. You can use cross-over Ethernet cables for two-node clusters.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.

- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.
- 13 Close the window.

Establishing a Microsoft failover cluster

Before installing SFW, you must first verify that Microsoft failover clustering is enabled (if a new installation of Windows Server 2008), and then establish a Microsoft failover cluster. This section summarizes the tasks; refer to Microsoft documentation for complete details.

To enable Microsoft failover clustering

- 1 In Server Manager, select **Features** in the left pane (tree view) and then click **Add Features** (link on the right side of the screen).
- 2 In the Add Features Wizard, check the **Failover Clustering** option, and click **Next**.
- 3 Click **Install**.
- 4 When the installation is complete, click **Close**.

To establish a Microsoft failover cluster

- 1 Ensure that you have met the hardware prerequisites for a failover cluster. You can run the Microsoft wizard to validate the configuration. See the Microsoft documentation for details.
- 2 Configure the shared storage and create a volume with drive letter “Q” for the cluster quorum disk. Use of other drive letters may result in the quorum recognition problems. You must have a basic disk reserved for this purpose on your shared storage.
Microsoft recommends a minimum of 500 MB for the quorum disk; refer to Microsoft documentation for specific requirements.
- 3 Create the first node of the cluster using Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 4 In the action pane, click **Create a Cluster**. The Create Cluster Wizard will start.
If this is the first time this wizard has been run, the Before You Begin page will appear. Review the information that is displayed and then click **Next**. You can hide this page on subsequent uses of the wizard; if this has been done, the first page might be the Select Servers page.
- 5 In the Select Servers panel, type the name of the first node in the Enter server name field and click **Add**. You can also use the Browse button to browse the Active Directory for the computers you want to add. Repeat this step for the second node.
- 6 After both nodes have been added to the list of Selected Servers, click **Next**.

- 7 Based on the information on the validation warning screen, assess your hardware configuration, and select one of the options. Symantec recommends that you select **Yes**, which starts the Validate a Configuration wizard. Follow the wizard instructions.
- 8 In the Access Point for Administering the Cluster screen, in the Cluster Name field, type the name for the failover cluster. This is the name that you use to connect to and administer the cluster.
- 9 In the Address field of the network area, type the appropriate IP address and then click **Next**.
- 10 In the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.
- 11 Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the 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, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so Microsoft clustering is now in control of the cluster storage on Server A, and both nodes of the storage cannot be accessed at the same time by the operating system.

To connect the two nodes

- 1 Connect corresponding cables between the three network cards on the two sites.
- 2 Connect the two switches at the two sites through the storage interconnect.
- 3 Test the connectivity between the two sites. Test the IP addresses of all the network adapter cards in the cluster. Bring up the command window and type `ping ipaddress`, where the *ipaddress* is the corresponding network adapter in the other node.

Installing SFW

This section assumes you are running a Microsoft failover cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 336.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 337.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 340.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups
See “[Moving the online groups](#)” on page 336.

Moving the online groups

If your resource groups are on the system where you are installing SFW, you must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.

If there is more than one resource group, you must repeat this step until all the resource groups are moved.

- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the Veritas Storage Foundation for Windows on a Microsoft failover cluster configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

Note: Before you install Storage Foundation for Windows, make sure that the node is inactive.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.
The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements.

The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).

The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.

- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.
 - Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.

If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.

To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.

Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/ Failover Cluster** option.

- Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
- To add a computer for installation, provide the name of the computer in the System Name box.
OR
If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.
 - To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
 - To know the verification status and other information of the added computer, click the information icon.
 - To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

■ **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

■ **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.
It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.
Click **Install** to install the software.
- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the online groups
See "[Moving the online groups](#)" on page 341.
- Completing the SFW installation
See "[Completing the SFW installation](#)" on page 341.

Moving the online groups

You can move the resource groups from the current system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the Microsoft failover cluster.

See “[SFW installation tasks](#)” on page 336.

Creating disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes on the cluster storage, which in a campus cluster consists of at least two storage arrays.

Before you create disk groups and volumes, consider the following items:

- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs which depend on the traffic load
- The disk groups and number of disks on each site

Note: For campus clusters, each disk group *must* contain an equal number of disks on each site.

- Types of volumes required and location of the plex of each volume in the storage array

Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Create two cluster disk groups on the storage—one for the application data files and one for the mirrored quorum.

For the Microsoft Exchange application data files, you could create a separate disk group for each storage group. It is best to separate data files from log files and place them in separate volumes. For example, you might create an Exchange disk group for the First Storage Group of the first Exchange virtual server (also called clustered mailbox server), EVS1_SG1_DG, with the following two volumes:

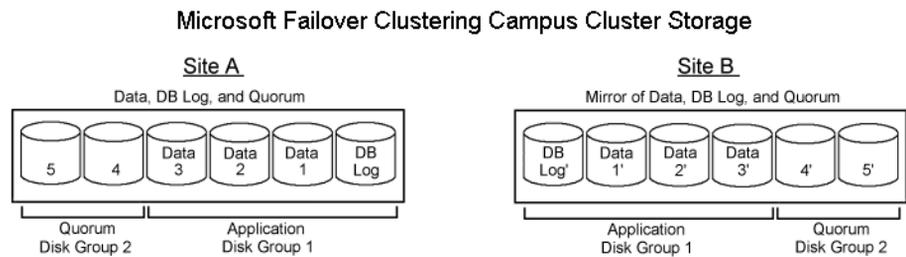
- EVS1_SG1_DB1: Contains the Exchange database. Each database in an Exchange storage group typically resides on a separate volume.
- EVS1_SG1_LOG: Contains the transaction log for the storage group.

The illustration that follows shows a typical Microsoft failover cluster with a campus cluster setup of disks. This example has only one application disk group that spans the storage arrays at both sites. The data and database log on Site A are mirrored to Site B. Each mirrored volume can have more than two disks, but must have an even number, such as four. All the application data could be in one large mirrored volume with multiple disks, but the same number of disks are

required on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

In the example, a four-way mirror for the quorum volume provides additional redundancy. The minimum configuration would be a two-way mirror. If possible, use small disks for the quorum volume. Microsoft recommends 500 MB for the quorum volume.

Figure 14-3 Microsoft failover cluster with campus cluster disks and disk groups example



Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site, and that each volume is a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 344
- [“Creating a volume”](#) on page 346

Considerations when creating new volumes

- For campus clusters, when you create a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.

- Symantec recommends using either simple mirrored (concatenated) or striped mirrored for the new volumes. Striped mirrored gives you better performance compared to concatenated.
When selecting striped mirrored, select two columns in order to stripe one enclosure that is mirrored to the second enclosure.
- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
The internal names for the disks which the current system can access for available storage are displayed, with names **Harddisk1**, **Harddisk2**, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

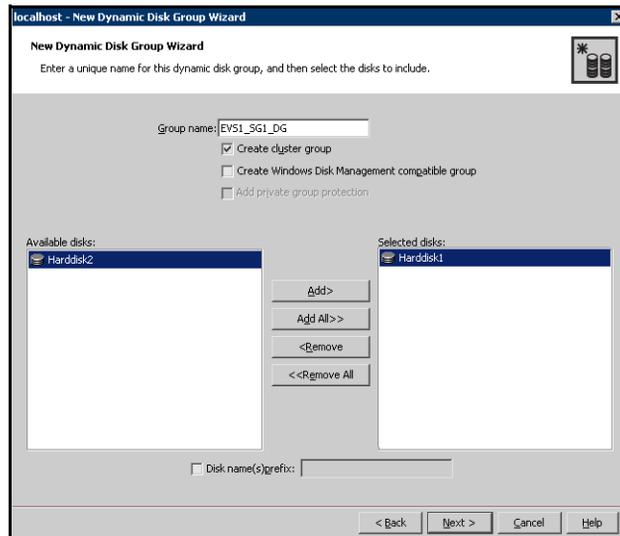
Use the following procedure to create a dynamic cluster disk group.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:



- Enter the name of the disk group (for example, EVS1_SG1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier.

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.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

■ Click **Next**.

- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Proceed to create the appropriate volumes on each disk.

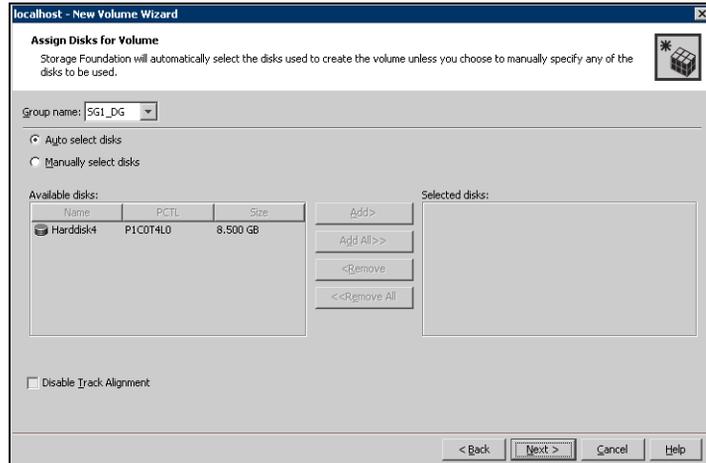
Creating a volume

Use the following procedure to create dynamic volumes.

To create dynamic volumes

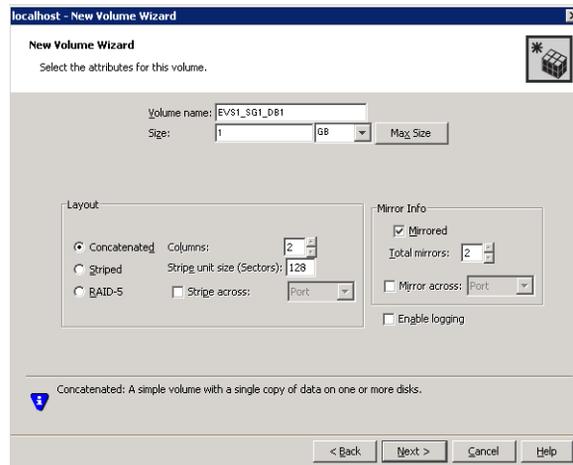
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example SG1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



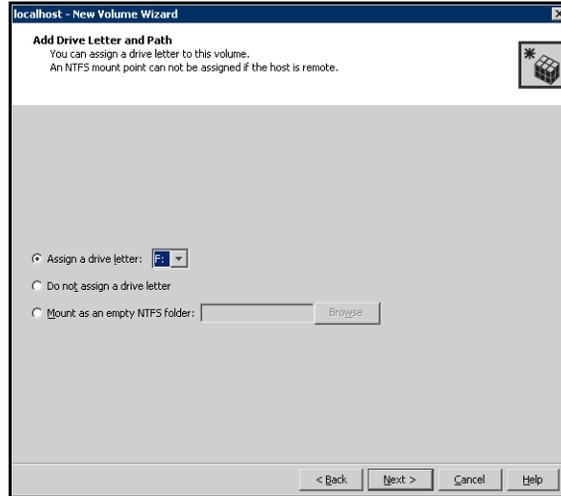
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the volume attributes.

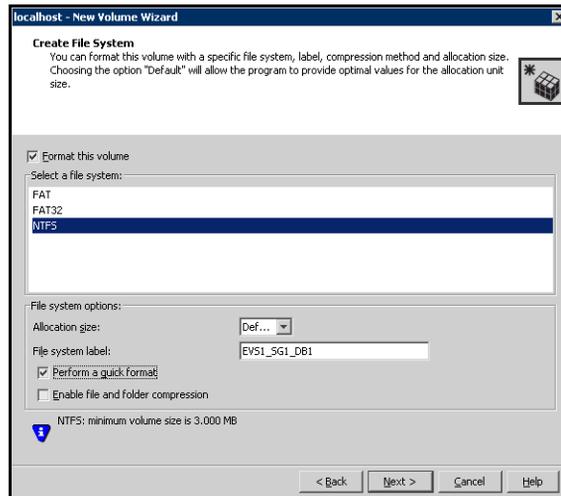


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation size or accept the Default.

- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.

13 Click **Finish** to create the new volume.

14 Repeat these steps to create the log volume (EVS1_SG1_LOG) in the EVS1_SG1_DG disk group. (The EVS1_SG1_LOG volume resides on its own disk.) For Exchange 2003, you could also create an MTA volume (EVS1_SG1_MTA).

15 If additional databases for EVS1_SG1_DG exist, create a volume for each database (for example, EVS1_SG1_DB2 and EVS1_SG1_DB3).

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Implementing a dynamic quorum resource

One of the key advantages of using SFW with Microsoft clustering is that you can create a mirrored quorum resource that adds fault tolerance to the quorum. The tasks for creating a mirrored quorum resource are:

- [“Creating a dynamic cluster disk group and a mirrored volume for the quorum resource”](#) on page 351
- [“Adding the Volume Manager Disk Group resource for the quorum”](#) on page 352
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 353

Creating a dynamic cluster disk group and a mirrored volume for the quorum resource

If you have not already completed this step, use SFW to create a cluster disk group for the quorum disks. Symantec recommends using four (small) disks for the mirrored quorum volume; you need a minimum of two disks. Microsoft recommends 500 MB for the quorum disk.

To create a four-way mirrored volume using the New Volume wizard

- 1 Create the cluster disk group with four small disks.
- 2 Create a volume with the four disks.
- 3 Select the **Concatenated** layout, click the **Mirrored** check box, and specify four mirrors.

For full details on a creating cluster disk groups and volumes, see [“Creating disk groups and volumes”](#) on page 342.

Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by Microsoft clustering.

Adding the Volume Manager Disk Group resource for the quorum

You must add the Volume Manager Disk Group resource corresponding to the disk group that you created for the quorum.

To add a Volume Manager Disk Group resource for the quorum in a Windows Server 2008 cluster

- 1 If Failover Cluster Management is already open, then proceed to Step 2. To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.
- 2 Verify that the cluster is online on the same node where you created the disk group.
- 3 In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
- 4 Right-click the new group and rename it, for example `QUORUM`.
- 5 Right-click `QUORUM` and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 6 Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
- 7 In the General tab of the Properties dialog box, type a name for the resource in the Resource Name field, for example, `QUORUM_DG_RES`.
- 8 On the Properties tab, in the Disk Group Name field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
- 9 Right-click the Quorum disk group resource (for example, `QUORUM_DG_RES`) in the left pane and select **Bring this resource online**.
The specified disk group resource, `QUORUM_DG_RES` resource, is created under the Quorum group (for example, `QUORUM`).

Changing the quorum resource to the dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

- 1 In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.
The Configure Cluster Quorum Wizard opens.
- 2 Review the screen and click **Next**.
- 3 Select either the **Node and Disk Majority** or **No Majority: Disk Only** radio button, and click **Next**.
- 4 Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.
This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, `QUORUM_DG_RES`.
- 5 Review the information in the Confirmation screen and click **Next**.
- 6 Click **Finish** to close the wizard.

Adding a Volume Manager Disk Group resource for Exchange 2007 installation

Before installing Exchange 2007, you add a Volume Manager Disk Group (VMDG) resource for the disk group that you created for the First Storage Group. By doing so, you can install the First Storage Group on a dynamic volume.

You add this resource to the Quorum group temporarily. After installation, you will move it to the Exchange group created by installation and set the appropriate dependencies.

Before creating this resource, start the cluster service on all the nodes in the cluster.

To create a Volume Manager Disk Group resource for the application

- 1 In the left pane (tree view) of the Failover Cluster Management tool, right-click the Quorum resource group under Services and Applications. Select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 2 In the General tab of the Properties dialog box, type a resource name for the new Volume Manager Disk Group (for example, EVS1_SG1_RES).
- 3 In the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the First Storage Group (for example, EVS1_SG1_DG).
- 4 Click **OK** to close the dialog box.
- 5 To bring the resource online, right-click the newly named resource and click **Bring this resource online**.

Installing the application on the cluster nodes

Exchange 2007 requires service pack 1 for Windows Server 2008.

Refer to the Microsoft documentation for prerequisites and details on installing a clustered mailbox server. You must install the Active Clustered Mailbox role on the first (active) node and install the Passive Clustered Mailbox role on the failover node.

Make sure that the same drive letter and path are available on all nodes and have adequate space for the installation. For example, if you install the Exchange program files in `C:\Program Files\ExchSrvr` on one node, you must install the files in `C:\Program Files\ExchSrvr` on all the other nodes. Use a drive letter on the local disk of each node in the cluster.

The First Storage Group database files must be installed on the shared storage. Browse to the dynamic volume that was previously prepared for the database. After installation, you can move the log to a separate volume.

Adding the Volume Manager Disk Group resources to the Exchange group

Exchange 2007 installation creates an application group with resources required by Exchange.

You need to move the Volume Manager Disk Group (VMDG) resource you added earlier for Exchange 2007 to the Exchange group.

In addition, add Volume Manager Disk Group resources for any other disk groups that you created for the Exchange application. Before creating this resource, start the cluster service on all the nodes in the cluster.

To move the Volume Manager Disk Group resource to the application group

- 1 In the quorum group, right click the resource you created for the Exchange First Storage Group and select the option to move it to another group.
- 2 Select the Exchange group as the target for the move and click **OK**.

To add Volume Manager Disk Group resources for the application

- 1 In the left pane (tree view) of the Failover Cluster Management tool, right-click the resource group under Services and Applications. Select **Add a resource > More resources > Add Volume Manager Disk Group**. A Volume Manager disk group resource is automatically created.
- 2 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** (in the center pane) to open the Properties dialog box. You can also right-click **New Volume Manager Disk Group** and select **Properties**.
- 3 In the General tab of the Properties dialog box, type the resource name for the New Volume Manager Disk Group (for example, ExchDG).
- 4 In the Properties tab, type the name of the disk group for which you want to add a resource.
- 5 Click **OK** to close the dialog box.
- 6 To bring the resource online, right-click the newly named disk group and click **Bring this resource online**.

Setting the database dependency on the disk group resource

After adding the Volume Manager Disk Group resources to the Exchange group, you must set the Exchange database resources to be dependent on them.

To set the database resource dependency on the VMDG resource

- 1 In Failover Cluster Management, select the Exchange resource group.
- 2 In the result pane, under Other Resources, right-click the appropriate database resource and select Properties.
- 3 In the Dependencies tab of the Properties dialog box:
 - Click the box **Click here to add a dependency**.
 - Select the appropriate Volume Manager Disk Group resource from the dropdown list of available resources.
 - Click **Insert**.
- 4 Click **OK** to close the Properties dialog box.
- 5 Repeat steps 2 through 4 for each additional database resource that exists in the Exchange group.

Moving the Exchange databases and logs to the dynamic volumes

During Exchange installation, the First Storage Group is installed to a dynamic volume. You must move the log to a separate volume.

If you created any other Exchange storage groups that were not located on the SFW dynamic volumes, move them to the dynamic volumes.

For each Exchange 2007 storage group, set the log files path and system files path to point to the log volume. For each database, set the database path to point to the appropriate database volume.

You can use either the Exchange Management Console or the Exchange Management Shell cmdlets to specify log and database locations. You can specify locations when creating new storage groups or databases or change the locations for existing storage groups or databases.

Note: You cannot use the Exchange Management Console to change the log file location for remote Mailbox servers.

The following procedures can be used to change paths for existing storage groups and databases.

See the Microsoft Exchange 2007 documentation for additional information on creating new storage groups and databases.

To use the Exchange Management Console to change log file and system file locations for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the storage group for which you want to change the log file location.
- 4 In the work pane, right-click the storage group for which you want to change the log file location and click **Move Storage Group Path**. The Move Storage Group Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the log files, and create a folder for the log files. Repeat for the system files. The volume should be the same for both.
- 6 Click **Move**. A warning appears that all databases in the storage group must be temporarily dismounted, which will make them inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the path was changed successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Storage Group Path wizard.

To use the Exchange Management Console to change the database file location for an existing database

- 1 Start the Exchange Management Console (**Start > All Programs > Microsoft Exchange Server 2007 > Exchange Management Console**) on the server on which the storage group is located.
- 2 In the console tree, expand **Server Configuration**, and then click **Mailbox**.
- 3 In the result pane, click the Mailbox server that contains the database for which you want to change the database file location.
- 4 In the work pane, right-click the desired database and click **Move Database Path**. The Move Database Path wizard appears.
- 5 On the Introduction panel, click **Browse**, specify the SFW volume to which to move the database file, and create a folder for it as needed.

- 6 Click **Move**. A warning appears indicating that the database must be temporarily dismounted, which will make it inaccessible to any user. To continue, click **Yes**.
- 7 On the Completion panel, confirm whether the database files were moved successfully. A status of Completed indicates that the wizard completed the task successfully. A status of Failed indicates that the task was not completed. If the task fails, review the summary for an explanation, and then click **Back** to make any configuration changes. Click **Finish** to complete the Move Database Path wizard.

Verifying the cluster configuration

After completing the configuration, verify that failover occurs as desired.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Failover Cluster Management to move all the resource groups from one node to another.
- Simulate a failover by shutting down an active cluster node.
Do not simulate failover in a production environment.

To move online groups

- 1 Open the Failover Cluster Management tool (**Start > All Programs > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

To shut down an active cluster node

- 1 Shut down the active cluster node normally.

- 2 Open Failover Cluster Management. Click **Start > All Programs > Administrative Tools > Failover Cluster Management** from any node in the cluster.
- 3 In Failover Cluster Management, verify that the owner name has changed. This confirms that all the resource groups have moved to another node.
- 4 If you need to move the resource groups back to the original node, restart the node you shut down in [step 1](#), select the resource group, and use **Move this service or application to another node > Move to node [name of node]** to move the resource group.

Deploying SFW and VVR with MSCS: New Exchange installation

This chapter includes the following topics:

- [Tasks in a new Exchange installation with SFW, VVR, and MSCS \(Windows Server 2003\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the primary site](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Setting up the secondary site \(Exchange 2003\)](#)
- [Setting up the secondary site \(Exchange 2007\)](#)
- [Installing Exchange on the secondary site \(Exchange 2007\)](#)
- [Setting up the Exchange group on the secondary site \(Exchange 2007\)](#)
- [Moving the Volume Manager Disk Group resources to the Exchange group \(Exchange 2007\)](#)
- [VVR components overview](#)
- [Creating resources for VVR \(primary and secondary sites\)](#)
- [Setting up the replicated data sets \(RDS\) for VVR](#)
- [Creating the RVG resource \(primary and secondary sites\)](#)
- [Normal operations and recovery procedures](#)

Tasks in a new Exchange installation with SFW, VVR, and MSCS (Windows Server 2003)

This section describes how to install and configure Storage Foundation for Windows and Veritas Volume Replicator (VVR) with MSCS and Exchange 2003 or 2007 on Windows Server 2003.

Exchange 2007 also runs on Windows Server 2008. For information on deploying on Windows Server 2008, see:

- [Chapter 16, “Deploying SFW and VVR with Microsoft failover clustering: New Exchange installation”](#) on page 409

After setting up a SFW environment with MSCS for Exchange on a primary site, you can create a secondary or “failover” site for replication. This section provides information on how to install and configure the Exchange components on the primary and secondary sites, with the intent of creating a parallel setup for the Exchange service group on both sites. This environment includes an active/passive configuration with one to one failover capability.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details about VVR.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 15-1 Tasks for MSCS for Exchange with VVR

Objective	Tasks
“ Reviewing the requirements ” on page 365	<ul style="list-style-type: none"> ■ Verifying hardware and software requirements
“ Reviewing the configuration ” on page 368	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster ■ Reviewing the benefits of a dynamic mirrored quorum
Part 1	<ul style="list-style-type: none"> ■ “Configuring the primary site” on page 371.
“ Configuring the storage hardware and network ” on page 190	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“ Establishing an MSCS cluster ” on page 191	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish an MSCS cluster

Table 15-1 Tasks for MSCS for Exchange with VVR (Continued)

Objective	Tasks
“Creating the MSDTC resource (Exchange 2003)” on page 192	<ul style="list-style-type: none"> ■ Creating the MSDTC resource for Exchange 2003
“Installing SFW with MSCS/Failover Cluster option” on page 371	<ul style="list-style-type: none"> ■ Verifying the driver signing options for Windows 2003 systems ■ Installing SFW (automatic installation) ■ Installing Cluster Option for Microsoft Cluster Service (MSCS) (manual option) ■ Restoring driver signing options for Windows 2003 systems ■ Configuring VxSAS
“Configuring SFW disk groups and volumes” on page 201	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes ■ Mounting and unmounting volumes
“Preparing the forest and domain (Exchange 2003)” on page 208	<ul style="list-style-type: none"> ■ Setting up the forest and domain prior to the Exchange installation
“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 208	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Microsoft Exchange Server installation ■ Repeating the procedure on additional nodes
“Creating an Exchange virtual server group (Exchange 2003)” on page 210	<ul style="list-style-type: none"> ■ Forming a cluster group that includes the IP Address, Network Name, SFW Disk Group, and Exchange2003 System Attendant (SA) resources
“Moving Exchange databases and logs to shared storage (Exchange 2003)” on page 221	<ul style="list-style-type: none"> ■ If desired, altering the paths for the Exchange 2003 databases and transaction logs
“Moving Exchange databases and logs to shared storage (Exchange 2007)” on page 223	<ul style="list-style-type: none"> ■ If desired, altering the paths for the Exchange 2007 databases and transaction logs
“Implementing a dynamic mirrored quorum resource” on page 225	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Creating the quorum resource for the Cluster Group ■ Changing the quorum resource to a dynamic mirrored quorum resource.

Table 15-1 Tasks for MSCS for Exchange with VVR (Continued)

Objective	Tasks
“Verifying the cluster configuration” on page 228	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node
Part 2	<p>“Setting up the secondary site (Exchange 2003)” on page 382. Repeating the primary site steps at the secondary site.</p>
	<p>“Setting up the secondary site (Exchange 2007)” on page 382. Repeating some of the primary site steps at the secondary site, with some different steps for Exchange installation and setup.</p>
Part 3	<p>Adding the components for VVR replication. See “VVR components overview” on page 387.</p>
“Creating resources for VVR (primary and secondary sites)” on page 387	<ul style="list-style-type: none"> ■ Creating an IP address for the Replicated Volume Group (RVG). ■ Creating a network name resource for the Replicated Volume Group (RVG) ■ Creating the Replicator Log volumes for VVR.
“Setting up the replicated data sets (RDS) for VVR” on page 389	<ul style="list-style-type: none"> ■ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites
“Creating the RVG resource (primary and secondary sites)” on page 400	<ul style="list-style-type: none"> ■ Using the Cluster Administrator to set up an RVG resource for VVR. ■ Setting up the Exchange dependency on the RVG resource
“Normal operations and recovery procedures” on page 404	<ul style="list-style-type: none"> ■ Operating the MSCS cluster with VVR and recovering from a disaster

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

Supported software

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW in a Microsoft cluster on Windows Server 2003:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
 Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
 For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2003 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 15-4 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)

Table 15-4 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2003 Server (Mailbox server role required)	<ul style="list-style-type: none"><li data-bbox="654 317 1160 369">■ Windows Server 2003 x64 Standard Edition or Enterprise Edition (SP2 required for all editions)<li data-bbox="654 378 1160 430">■ Windows Server 2003 R2 x64 Standard Edition, Enterprise Edition (SP2 required for all editions)

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. The following table summarizes disk space requirements for SFW.

Table 15-5 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

Observe the following system requirements:

- The configuration described requires shared disks to support applications that migrate between nodes in each cluster.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- MSCS requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. There are six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.
- Each system requires 1 GB of RAM for SFW.
- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- A minimum 256 MB of RAM per server is required for Exchange 2003; refer to your Microsoft documentation for more information.

- Systems to be clustered must be configured as part of a Windows Server 2003 domain. Each system in an MSCS cluster must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and MSCS software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- The administrator account for the Exchange virtual server on the primary site must be the same account used for the Exchange virtual server on the secondary site.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft Exchange documentation for instructions on creating a reverse lookup zone.
- For Exchange 2003, you must be an Enterprise Administrator, Schema Administrator, Domain Administrator, and Local Machine Administrator to run ForestPrep. You must be a Domain Administrator and Local Machine Administrator to run DomainPrep. Refer to the Microsoft documentation for permissions requirements.
- Symantec recommends two disks for Exchange: one for Exchange database files and one for Exchange log files.
- The cluster on the secondary site must reside in the Active Directory domain of the cluster on the primary site.

Note: Refer to the Hardware Compliance List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

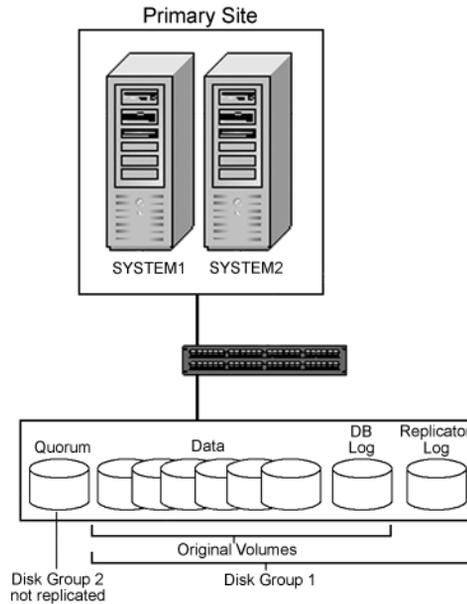
Reviewing the configuration

This overview highlights the high availability within a cluster, and a clustered configuration with VVR between two sites. In an active/passive configuration, one or more Exchange virtual servers can exist in a cluster. In this case, Exchange virtual server EVS1 can fail over from SYSTEM1 to SYSTEM2 on the primary site, and SYSTEM5 to SYSTEM6 on the secondary site.

Each site has a cluster with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Exchange application data is stored on the volumes that are under the control of the RVG. A separate disk group is created for the quorum volume which is not replicated.

[Figure 15-1](#) shows a cluster configuration on the primary site.

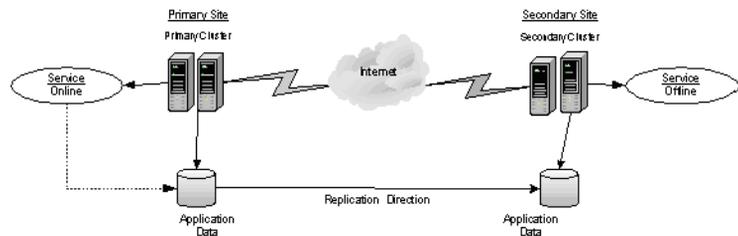
Figure 15-1 DR configuration on the primary site



In a two-site VVR environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. The following illustration displays an environment that is prepared for a disaster with a two-site VVR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 15-2 shows an example disaster recovery configuration before a failure.

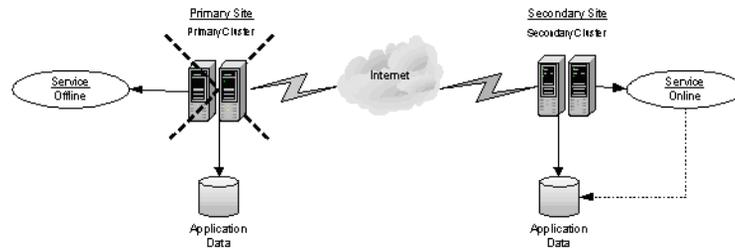
Figure 15-2 Disaster recovery example configuration



When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the data that was replicated to the secondary site can be used to restore the application services to clients.

Figure 15-3 shows an example disaster recovery configuration after a failure.

Figure 15-3 Disaster recovery configuration after a failure



For a VVR configuration with MSCS, the following apply:

- An MSCS cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.

In an MSCS cluster without SFW, the quorum disk is a point of failure because MSCS only supports a basic physical disk and does not enable you to mirror the quorum resource.

The main advantage of SFW is that it provides a dynamic mirrored quorum resource for MSCS. If the quorum resource fails, the mirror takes over for the resource. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose. You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that Exchange is working in the cluster before adding the dynamic quorum volume.

- SFW enables you to add fault-tolerance to data volumes. Symantec recommends mirroring log volumes and a mirrored striped RAID layout for data volumes. SFW offers multiple disk groups, mirrors, capacity management and automatic volume growth, online storage migration, performance tuning, hot relocation, dirty region logging, RAID-5 logging, dynamic multi-pathing, and enhanced snapshot capabilities with FlashSnap. Refer to the *Veritas Storage Foundation Administrator's Guide* for details on these features.

Configuring the primary site

The beginning steps for configuring the primary site for a disaster recovery solution are the same as when configuring for high availability.

You begin with the following topics in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):

- [“Configuring the storage hardware and network”](#) on page 190
- [“Preparing the forest and domain \(Exchange 2003\)”](#) on page 208
- [“Establishing an MSCS cluster”](#) on page 191
- [“Creating the MSDTC resource \(Exchange 2003\)”](#) on page 192

You can then continue with the following procedure for installing SFW for a disaster recovery configuration.

In configuring the primary site to prepare for a disaster recovery, you need to make sure that the VVR option is selected during the SFW installation. After installation you must configure the Veritas Volume Replicator Security Service (VxSAS).

Installing SFW with MSCS/Failover Cluster option

This section assumes you are running a Microsoft cluster and you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft cluster simultaneously.

SFW installation tasks

The product installer enables you to install the software for Veritas Storage Foundation 5.1 for Windows. The installer automatically installs SFW. You must select the options to install VVR, and the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation. The steps in this section are based on a server installation.

Installing SFW involves the following:

- Performing pre-installation tasks
See “[Pre-installation tasks](#)” on page 372.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 374.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 377.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
See “[Changing the driver signing options](#)” on page 372.
- Moving the Online Groups
See “[Moving the online groups](#)” on page 373.

Changing the driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

Note: The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft. This option installs a Symantec Trusted certificate on the systems you select for installation. If this option is selected, you do not need to set the driver signing options to Warn or Ignore.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 15-6 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

To change the driver signing options on each system

- 1 Log on locally to the system.
 - 2 Open the Control Panel and click **System**.
 - 3 Click the **Hardware** tab and click **Driver Signing**.
 - 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
 - 5 Click **OK**.
 - 6 Repeat for each computer.
- If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Moving the online groups

If your resource groups are online on the system where you are installing SFW. You must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open the Cluster Administrator (**Start > All Programs > Administrative Tools > Cluster Administrator**).
- 2 Right-click on a resource group and click **Move Group**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Cluster Administrator console, verify that the owner name has changed. This confirms that all the resource groups have moved to another system.
- 4 If you need to move all the resource groups back to the original system use **Move Group** to move all the resource groups.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install Veritas Storage Foundation for Windows on a Microsoft cluster configuration.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.
The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements. The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).
The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.
- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.

- Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the Complete/Custom link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.
 If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.
 To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.
 Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Volume Replicator (VVR)** option.
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** option.
 - Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
 - To add a computer for installation, provide the name of the computer in the System Name box.
 OR
 If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.

- To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
- To know the verification status and other information of the added computer, click the information icon.
- To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

■ **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

■ **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install

Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.

It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.

Click **Install** to install the software.

- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups
See “[Moving the online groups](#)” on page 377.
- Completing the SFW Installation for the other systems in the MSCS cluster
See “[Completing the SFW installation](#)” on page 378.
- Configure the VxSAS service
See “[Configuring the VxSAS service](#)” on page 378.
- Resetting the driver signing options
See “[Resetting the driver signing options](#)” on page 381.

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open Cluster Administrator (**Start>Control Panel>Administrative Tools>Cluster Administrator**). Connect to the appropriate cluster through the console.

- 2 From the configuration tree, right-click **Groups**, and click **Move Group**.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 372.

Configuring the VxSAS service

Complete the following procedure to configure this required service for VVR. This procedure should not be done until you have installed SFW on all cluster systems. Otherwise, you will get an error message from the VxSAS wizard if you try to select a system without SFW installed.

You can run the VxSAS wizard from the secondary site once SFW is installed on all cluster systems; at that time, you can run the wizard for both the primary and secondary site systems. The MSCS groups can be either online or offline.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. Accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

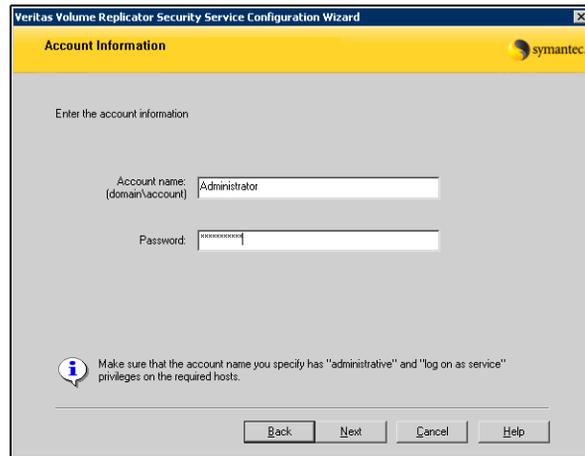
- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxSASCfg.exe` from the command prompt of the required machine.

The welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

2 Complete the Account Information wizard page as follows:

- | | |
|----------------------------------|--|
| Account name
(domain\account) | Enter the administrative account name in the Account name field. |
| Password | Specify a password in the Password field. |

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same user name and password when configuring the VxSAS service on the other hosts.



3 After providing the required information, click **Next**.

4 Select the required domain to which the hosts that you want to configure belong, from the Domain Selection wizard page.

Selecting Domains The Available Domains pane lists all the domains that are present in the Windows network neighborhood. Select the required domain by moving the appropriate name from the Available Domains pane to the Selected Domains pane, either by double-clicking it or using the arrow button.

Adding a Domain If the domain name that you require is not displayed, then add it by using the **Add Domain** option. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected Domains list.

After specifying the domain, click **Next**.

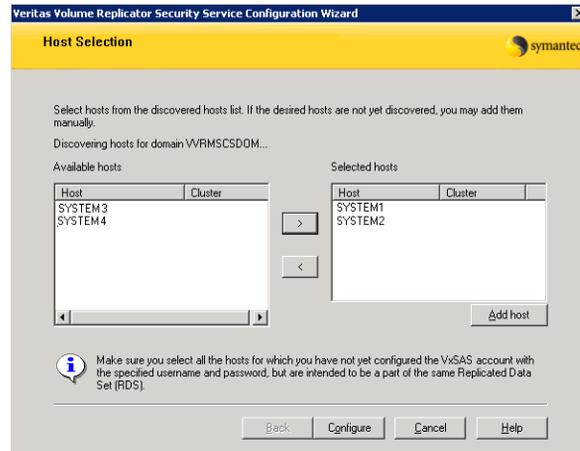
- 5 Select the required hosts from the Host Selection page.

Selecting Hosts The Available Hosts pane lists the hosts that are present in the specified domain.

Select the required host by moving the appropriate name from the Available Hosts list to the Selected Hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a Host If the host name you require is not displayed, then add it using **Add Host** option. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected Hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.



- 6 After the configuration completes, the Configuration Results page is displayed. If the operation is successful, then the Status column displays the appropriate message to indicate that the operation was successful.

If the operation was not successful, then the page displays the status as failed and the corresponding details on why the account update failed. It also displays the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

- 7 Click **Finish** to exit the wizard.

Resetting the driver signing options

You must reset the driver signing options to its previous state. This is to ensure a secure system environment.

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

After completing SFW installation on the primary site, continue with the additional tasks for completing the primary site configuration.

Completing the primary site configuration

After you have installed SFW, the remaining steps for configuring the primary site for disaster recovery are the same as when configuring high availability.

To complete the primary site configuration, refer to the following topics as required for Exchange 2003 or Exchange 2007, in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):

- [“Configuring SFW disk groups and volumes”](#) on page 201
- [“Adding a Volume Manager Disk Group resource for Exchange 2007 installation”](#) on page 208
- [“Installing Exchange Server”](#) on page 209
- [“Creating an Exchange virtual server group \(Exchange 2003\)”](#) on page 210
- [“Adding Volume Manager Disk Group resources to the Exchange 2007 group”](#) on page 220
- [“Moving Exchange databases and logs to shared storage \(Exchange 2003\)”](#) on page 221
- [“Moving Exchange databases and logs to shared storage \(Exchange 2007\)”](#) on page 223
- [“Implementing a dynamic mirrored quorum resource”](#) on page 225

- [“Verifying the cluster configuration”](#) on page 228

Setting up the secondary site (Exchange 2003)

Setting up the secondary site works somewhat differently for Exchange 2003 and Exchange 2007. This topic covers the Exchange 2003 environment.

After setting up a SFW environment with MSCS and Exchange 2003 on the primary site, complete the same tasks on the secondary site prior to the Exchange installation; this assumes you are already aware of the material in [“Reviewing the requirements”](#) on page 365.

See the following topics for the sequence of tasks:

- [“Configuring the primary site”](#) on page 371
- [“Completing the primary site configuration”](#) on page 381

Note the following special considerations when setting up the secondary site for Exchange 2003:

- During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume names and sizes
 - Drive letters
- Make sure to take the Exchange service group offline on the primary site; otherwise, the installation on the secondary site will not function properly.
- Specify the same name for the Exchange virtual server on the Primary and Secondary sites.

After completing these tasks you will have a clustered secondary site with SFW with the MSCS option and Exchange installed on all the nodes on both the Primary and the Secondary Sites.

Setting up the secondary site (Exchange 2007)

After setting up an SFW environment with Microsoft failover clustering and Exchange 2007 on the primary site, configure the secondary site. Many tasks are the same as on the primary site.

The following is the list of tasks in the recommended order:

- Configure the storage hardware and network and establish a Microsoft failover cluster.

Refer to the following topics in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):

- [“Configuring the storage hardware and network”](#) on page 190
- [“Establishing an MSCS cluster”](#) on page 191
- Install SFW for a disaster recovery configuration and configure the VxSAS service. Refer to the following topic in this chapter:
 - [“Installing SFW with MSCS/Failover Cluster option”](#) on page 371
- Configure SFW disk groups and volumes.
 During the creation of disk groups and volumes for the secondary site, make sure the following are exactly the same as the cluster on the primary site:
 - Cluster disk group names
 - Volume names and sizes
 - Drive letters

Refer to the following topic in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):

- [“Configuring SFW disk groups and volumes”](#) on page 201
- Add Volume Manager Disk Group resources for the Exchange 2007 disk group(s)
 The resource will not be used during the Exchange 2007 installation but will be used in the process of recovering the existing database and log information from Active Directory after Exchange installation.
 If you created any other Exchange storage groups, include resources for those disk groups as well.
 Refer to the following topic in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):
 - [“Adding a Volume Manager Disk Group resource for Exchange 2007 installation”](#) on page 208

- Install Exchange 2007
The procedure for installing and configuring Exchange 2007 on the secondary site is different from on the primary site. See the following topic in this chapter:
 - [“Installing Exchange on the secondary site \(Exchange 2007\)”](#) on page 384
 - Set up the Exchange group on the secondary site
Unlike on the primary site, Exchange installation does not create the Exchange group and resources. However, after Exchange installation, you can run a command that recovers the resource information from Active Directory. See the following topic in this chapter:
 - [“Setting up the Exchange group on the secondary site \(Exchange 2007\)”](#) on page 385
 - Move the Volume Manager Disk Group resources to the Exchange group
Move the disk resources you added for Exchange from the cluster group to the Exchange group. See the following topic in this chapter:
 - [“Moving the Volume Manager Disk Group resources to the Exchange group \(Exchange 2007\)”](#) on page 386
 - Implement a dynamic mirrored quorum resource
Refer to the following topic in [Chapter 11, “Deploying SFW with MSCS: New Exchange installation”](#):
 - [“Implementing a dynamic mirrored quorum resource”](#) on page 225
- After completing these tasks you will have a clustered site with SFW with the Microsoft cluster option and Exchange installed on all the nodes on both the primary and the secondary sites.

Installing Exchange on the secondary site (Exchange 2007)

Review the following requirements for installing Exchange Server 2007 on the secondary site:

- Before you begin installation, make sure to take the Exchange cluster group offline on the primary site; otherwise, the installation on the secondary site will not function properly.
- On the secondary site nodes, do not install the Active Clustered Mailbox role. Instead, install the Passive Clustered Mailbox role on all nodes of the secondary site.

- Make sure that the same drive letter and path are available on all nodes and have adequate space for the installation. For example, if you install the Exchange program files in C:\Program Files\ExchSrvr on one node, you must install the files in C:\Program Files\ExchSrvr on all the other nodes. Use a drive letter on the local disk of each node in the cluster.

Refer to the Microsoft documentation for additional prerequisites and details on Exchange installation.

Setting up the Exchange group on the secondary site (Exchange 2007)

Installing Exchange 2007 on the secondary site does not create the Exchange group and resources as it does on the primary site. However, after Exchange installation, you can use the RecoverCMS command to recover the Exchange group and resource information from Active Directory.

To set up the Exchange group on the secondary site

- 1 On the secondary site, ensure that you have set up the cluster disk groups and dynamic volumes to match those on the primary site.
- 2 Ensure that you have created the Volume Manager Disk Group resources that correspond with the Exchange disk groups.
- 3 Ensure that the disk groups are online and volumes mounted on the node where you are going to run the RecoverCMS command.
- 4 Change to the directory where you installed Exchange.
- 5 Run the RecoverCMS command, using the following syntax:

```
Setup.com /recoverCMS  
/CMSName:<name>/CMSIPAddress:<ip>
```

Where *<name>* is the name you assigned to the Exchange server during installation on the primary site and *<ip>* is a new IP, not the one assigned during Exchange installation on the primary site.
- 6 On the secondary site, bring the Exchange cluster group and its resources offline.
- 7 On the primary site, bring the Exchange cluster group and its resources online.

Moving the Volume Manager Disk Group resources to the Exchange group (Exchange 2007)

Before Exchange 2007 installation, you added Volume Manager Disk Group (VMDG) resources for Exchange to the Cluster Group so that they could be used for setting up Exchange on the secondary site. You now move the resources to the Exchange group.

Use the Cluster Administrator to move the Volume Manager Disk Group resources that you added for Exchange 2007 from the Cluster Group to the Exchange group. After doing so, set the Exchange database resource for the First Storage Group to depend on the Volume Manager Disk Group resource.

If you have created additional cluster disk groups for Exchange storage groups, add a Volume Manager Disk Group resource for each cluster disk group to the Exchange group. Set the appropriate dependencies.

VVR components overview

The next set of tasks is to configure both sites for VVR replication. You configure the following Veritas Volume Replicator components:

Replicated Volume Group (RVG)	<p>An RVG is made up of one or more volumes in a SFW disk group. The updates made on the RVG on the primary host are sent to a configured secondary host. Thus, on the secondary host 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. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG.</p> <p>An RVG within a disk group is the container for replication, so if you have multiple disk groups, you will 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.</p>
Replicated Data Set (RDS)	<p>An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).</p>
Replicator Log volume	<p>Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any RVG updates 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 needs to become the new primary site. The log volumes at the two sites must have the same name. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.</p>

Creating resources for VVR (primary and secondary sites)

Create the resources for VVR at the primary and secondary sites.

Create an IP Address resource and a Network Name resource for the RVG on both the primary and secondary site.

A separate valid IP address is necessary for VVR replication, because on the secondary cluster before a disaster, the EVS IP must be offline whereas the VVR IP must be online.

To create an IP Address resource for RVG

- 1 Right click on the created EVS group and select **New>Resource**. The Resource creation wizard appears.
- 2 Specify a name for the **IP Address** resource in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
Select the **IP address** from the **Resource Type** field drop-down list. Click **Next**. The **Possible Owners** page appears.
- 3 By default, all the nodes in the cluster are listed as possible owners. Click **Next**. The **Dependencies Page** appears.
- 4 Make sure the **Resource Dependencies** pane is empty, and click **Next**.
- 5 On the **TCP/IP Address Parameters** page, enter the IP address and corresponding subnet mask. This IP address should be used for replication while configuring VVR later on. Click **Finish** to create the **IP Address** resource.

To create a Network Name resource for Replicated Volume Group (RVG)

- 1 Right-click on the EVS group and select **New > Resource**. The Resource creation wizard appears.
- 2 Specify a name for the **Network Name** resource in the **Name** field.
If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Network Name** from the **Resource Type** field drop-down list. Click **Next**. The **Possible Owners** page appears.
- 3 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.
- 4 On the **Dependencies** page, select the IP Address resource that you just created for the RVG from the **Available Resources** pane and add it to the **Resource Dependencies** pane. Click **Next**. The **Parameters Page** appears.
- 5 In the **Name** field on the **Network Name Parameters** page, specify any name except the node and Exchange Virtual Server names. Click **Finish**.
The network name for the RVG must be different for the Primary and Secondary cluster.

Repeat the same procedure to create the IP and the Network Name resource at the secondary site. Bring the resources online.

Setting up the replicated data sets (RDS) for VVR

For each disk group you created for the application, you set up a Replicated Data Set (RDS) on the primary and secondary hosts. The Setup Replicated Data Set Wizard enables you to configure an RDS for both sites.

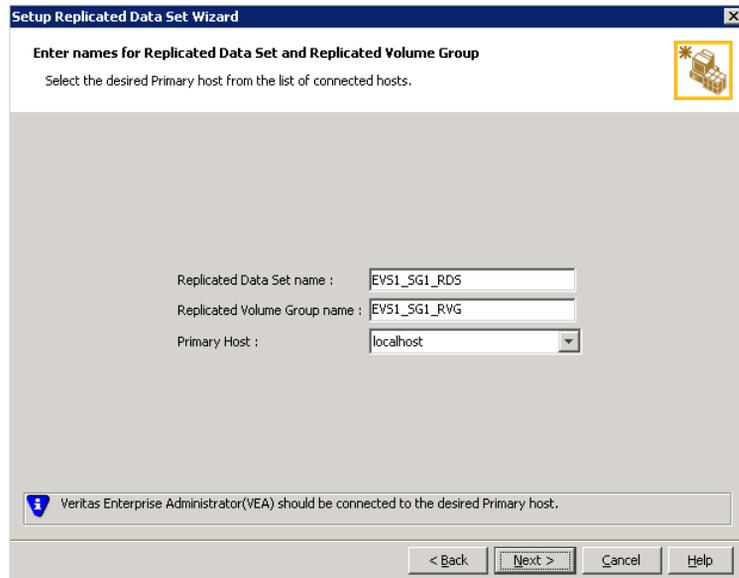
- Verify that the set of volumes that are intended to store the user Exchange database files, the user Exchange log files, and the Replicator Log file have been created on the primary node.
- Verify that the EVS IP resource is offline on the secondary site. This would also offline all the dependent Exchange resources.
- Verify that the data and Replicator Log volumes are not of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volume names containing a comma
- Verify that the Replicator Log volume does not have a DCM.
- Verify that the cluster disk group is imported on the primary and secondary site.
- Verify that VxSAS has been configured.

The following procedure enables you to set up an RDS on the primary and secondary sites and start replication.

To create the Replicated Data Set

- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Setup Replicated Data Set**.
- 3 Read the information on the Welcome page and then click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG) and then click **Next**.

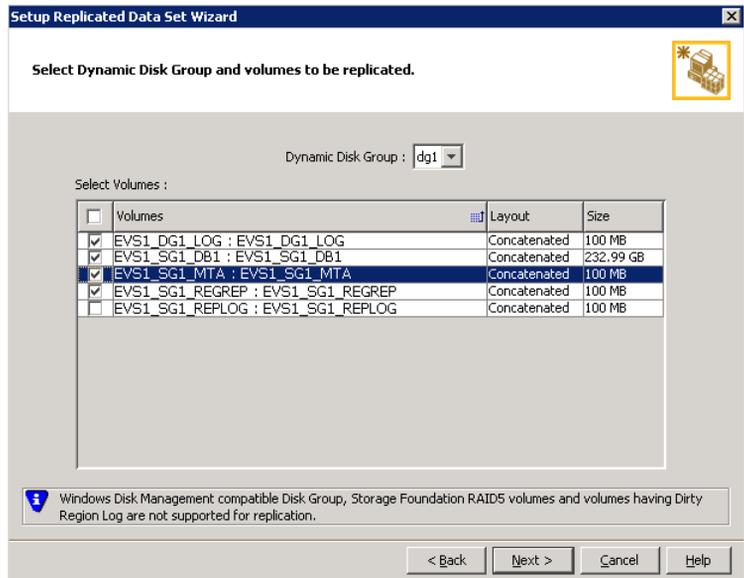


The screenshot shows a Windows-style dialog box titled "Setup Replicated Data Set Wizard". The main heading is "Enter names for Replicated Data Set and Replicated Volume Group". Below this, a sub-heading reads "Select the desired Primary host from the list of connected hosts." The dialog contains three input fields: "Replicated Data Set name" with the text "EV51_SG1_RDS", "Replicated Volume Group name" with the text "EV51_SG1_RVG", and "Primary Host" with a dropdown menu currently showing "localhost". At the bottom, there is an information icon and a message: "Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host." The bottom right corner features four buttons: "< Back", "Next >", "Cancel", and "Help".

By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

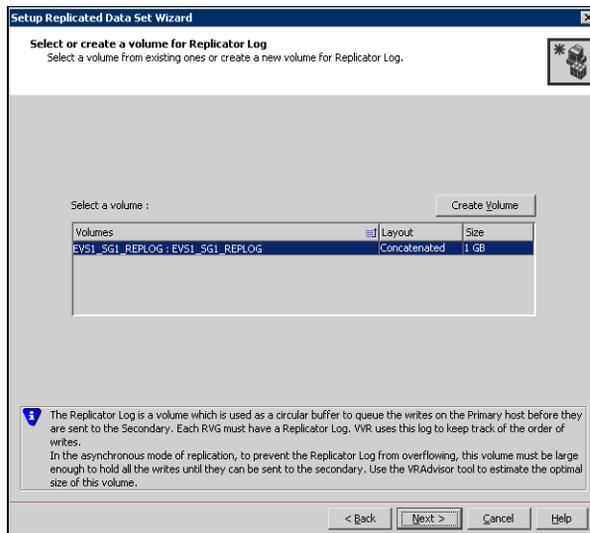
- 5 Select from the table the dynamic disk group and data volumes that will undergo replication and then click **Next**.



To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

6 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (EVSI_SG1_REPLOG).
If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

- Name** Enter the name for the volume in the **Name** field.
- Size** Enter a size for the volume in the **Size** field.
- Layout** Select the desired volume layout.

Disk Selection

Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** check box to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The check box will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this check box along with **Select disks manually** option, then the user can select only TP disks from **Available Disks**.

For more information on Thin Provisioning refer to the *Veritas Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want VVR to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on it or select **Add** to move the disks into the Selected disks pane.

- Click **OK** to create the Replicator Log volume.
- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.

7 Review the information on the summary page and click **Create Primary RVG**.

8 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

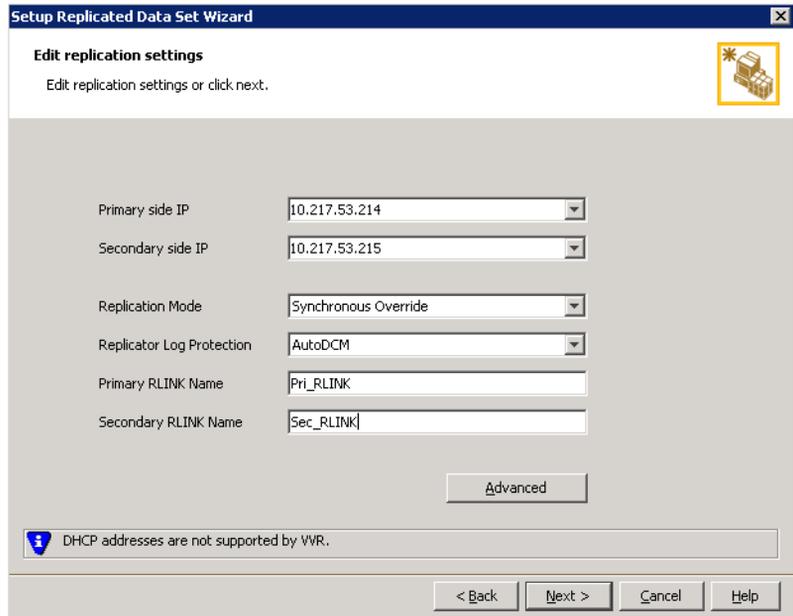
9 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field and then click **Next**.

If the Secondary host is not connected to VEA, the wizard tries to connect it when you click Next. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 10 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.
The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:
 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary
Otherwise, the RDS setup wizard enables you to create the required volumes manually.
 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
- 11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.
This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.
 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
 - If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.
Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.
When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.
- If all the data volumes to be replicated meet the requirements, this screen does not occur.

12 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:



- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Secondary side IP Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode	<p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as <code>MISSING</code>.</p>
Replicator Log Protection	<p>The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.</p> <p>The DCM option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.</p> <p>The Off option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to Off, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.</p>

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK
Name

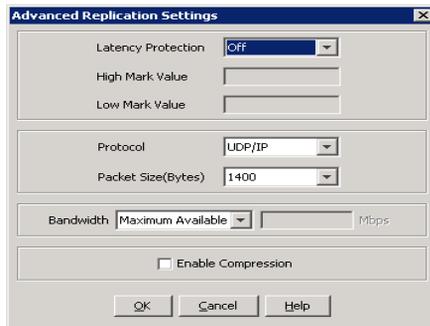
This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Secondary RLINK
Name

This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

- Click **Next** to start replication with the default settings.

- 13 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

Off is the default option and disables latency protection.

Fail enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.

Override enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Enable Compression Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box and then click **Next**.

14 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from Checkpoint If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu.
- Click **Next** to display the Summary page.

15 Review the information.

Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the RVG resource (primary and secondary sites)

Once the replicated data set (RDS) is configured for VVR, you create an RVG resource in the Exchange Virtual Server group and specify its dependencies.

To create a Replicated Volume Group (RVG) resource

- 1 Right click on the Exchange Virtual Server (EVS) group that you have created and select **New > Resource**. The **New Resource** wizard appears.
- 2 Specify a name for the **Replicated Volume Group** resource in the **Name** field. If required, you can add a description about the resource in the **Description** field.
Specify the resource type by selecting **Replicated Volume Group** from the **Resource Type** field drop-down list. Click **Next**. The **Possible Owners** page appears.
Configure a separate resource monitor process for the RVG resource by selecting the **Run this resource in a separate Resource Monitor** checkbox provided in the **New Resource** wizard.

- 3 By default all the nodes in the cluster are listed as possible owners. Click **Next**. The **Dependencies** page appears.
- 4 On the dependencies page, select the VVR Network Name resource, and the Disk Group resource, from the **Available Resources** and add it to **Resource Dependencies**. Click **Next**.
- 5 On the **Replicated Volume Group Parameters** page select the created RVG. Click **Finish**.
- 6 Repeat the same steps to create the RVG resource at the Secondary site.

Setting the System Attendant resource dependency on the RVG resource (Exchange 2003)

For Exchange 2003, the final step in cluster configuration is to set the System Attendant resource dependency on the RVG resource.

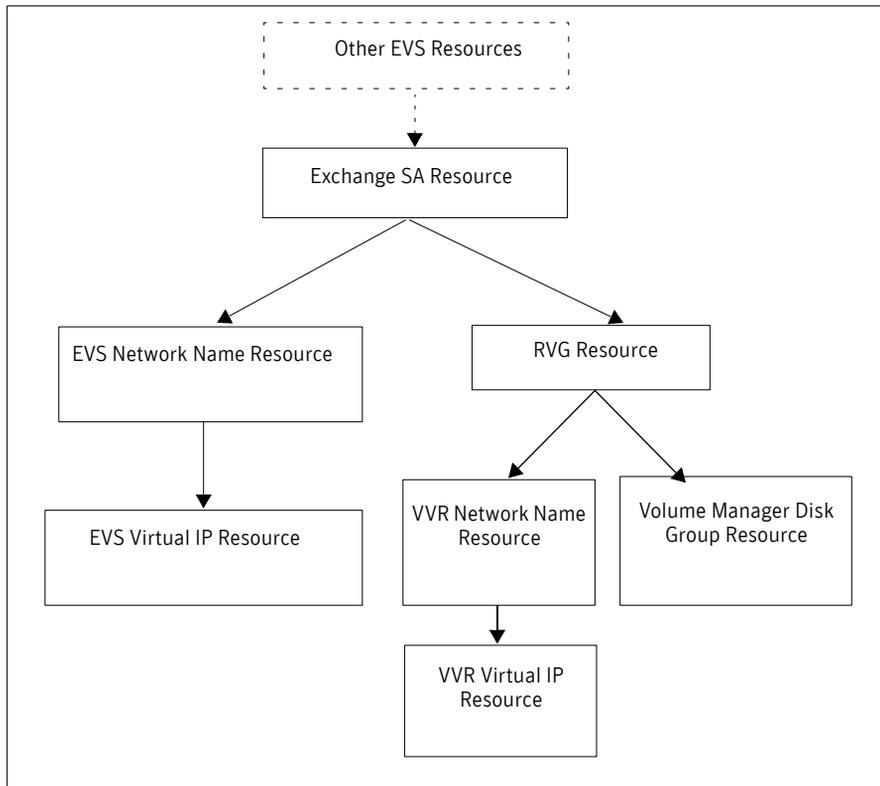
To set the System Attendant resource dependency on the RVG resource

- 1 Make sure the **Exchange System Attendant** resource is offline before attempting to modify the dependencies.
- 2 Right-click on the **Exchange System Attendant** resource and select **Properties>Dependencies** tab. This will display the Dependencies page.
- 3 Click **Modify**. Select the **Replicated Volume Group** resource from the **Available Resources** and add it to **Resource Dependencies**. Remove the **Disk Group** resource from **Resource Dependencies**. Click **OK**.

The cluster configuration is now complete. Online the entire EVS group on the Primary cluster. On the Secondary cluster only the **Disk Group** resource, VVR IP Address resource, VVR Network Name resource and the RVG resource should be online.

[Figure 15-4](#) shows a dependency graph with the dependencies that have been established.

Figure 15-4 Example Exchange 2003 dependency graph with an RVG



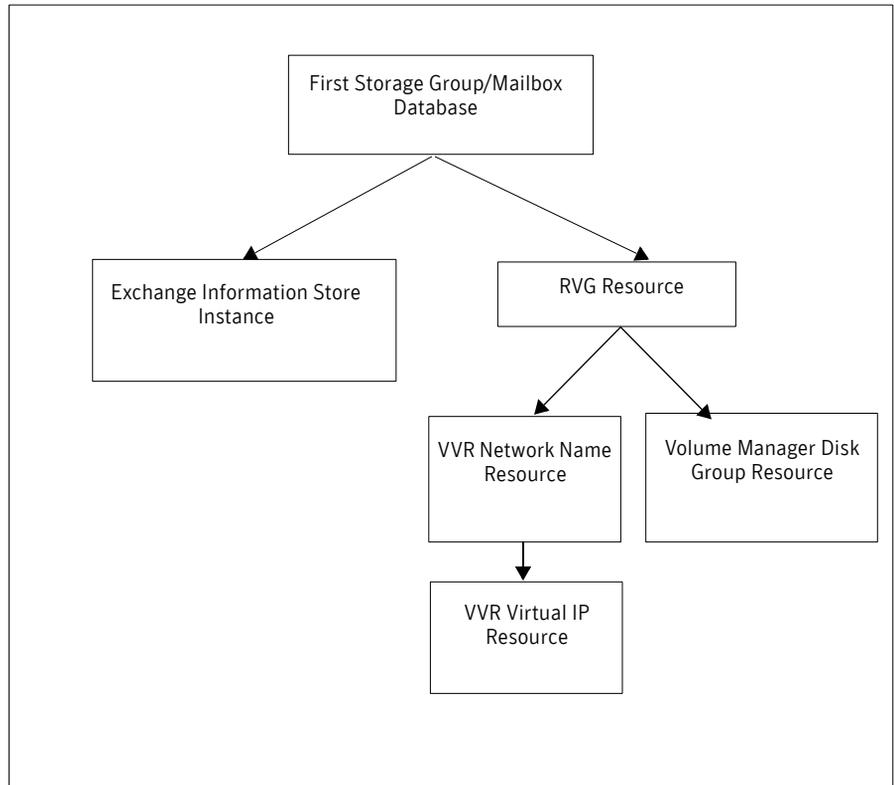
Setting the database resource dependency on the RVG resource (Exchange 2007)

For Exchange 2007, the Exchange database resource was earlier set to depend on a Volume Manager Disk Group resource that corresponded to the disk group created for the application. After you add the RVG resource for that disk group, you must change the dependency. You set the database resource to depend on the RVG resource, removing the dependency on the Volume Manager Disk Group resource.

The cluster configuration is now complete. Online the entire EVS group on the Primary cluster. On the Secondary cluster only the **Disk Group** resource, VVR IP Address resource, VVR Network Name resource and the RVG resource should be online.

Figure 15-5 shows a dependency graph with the dependencies that have been established for the RVG resource.

Figure 15-5 Example Exchange 2007 dependency graph with an RVG



Normal operations and recovery procedures

This section gives considerations for normal VVR operations and also describes the recovery process.

Normal operations

Under normal operating conditions you can monitor the status of the replication using the following:

- VEA GUI
- Command Line Interface (CLI)
- Performance Monitor (perfmon)
- Alerts

For details, refer to the “Monitoring Replication” section in the *Veritas Volume Replicator Administrator’s Guide*.

Performing planned migration

For maintenance purposes, or for testing the readiness of the Secondary host you may want to migrate the application to the Secondary host. These are a generic set of tasks that you may need to perform.

To migrate the application to the Secondary host

- 1 Take the RVG resource offline on both the clusters, which will take all dependent resources (such as the Exchange resources) offline.
- 2 Transfer the Primary role to the Secondary using the Migrate option. From the **VEA** screen, right-click the Primary RVG and select **Migrate**. Select the Secondary host and click **OK**. The replication role is migrated to the Secondary host.
- 3 Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as those of the original Primary.
- 4 Bring the RVG resource online on the new Secondary.
- 5 Bring the EVS group online on the new Primary.

You can now verify that the Exchange runs fine on the new Primary with the replicated data. After verifying you can revert back the roles to its original state using the same set of tasks described above.

Any changes that you make to the data on the new Primary will get replicated to the original Primary, which is now the Secondary.

Replication recovery procedures

This section provides information on bringing up an Exchange server on the Secondary Site, in the event of a disaster. It also explains how to migrate the Primary role back to the original Primary Site once it is in a good state after a disaster.

Bringing up Exchange on the secondary site

To bring up Exchange on the Secondary Site, recover the Exchange data.

To recover the Exchange data

- 1 From the left-pane in the VEA GUI console on a system on the Secondary Site, right-click on the desired secondary RVG node inside the replication network. Select the **Take Over** option. The **Take Over** dialog box is displayed.
 - By default, the **Enable Fast-Failback Logging** option is selected if the data volumes have DCM logs associated with them. You can use this option to perform takeover with fast-failback logging.

The DCM is activated for fast-failback logging and the new incoming writes are marked on the DCM of the New Primary.

If the replication status of Secondary RVG was *Inactive* when the Primary failed, then the **Enable Fast-Failback Logging** option is unavailable for selection. In this case you can perform **Take Over** without using fast-failback logging.
 - Select the **Synchronize Automatically** option if you want the new Primary and the original Primary to get synchronized automatically, after the original Primary recovers.

If you have not selected this option, the original Primary, after it recovers will be in the *Acting as Secondary* state. To synchronize this original Primary with the new Primary use the **Resynchronize Secondaries** option from new Primary RVG's right-click menu. When the resynchronization starts, the original Primary which was in the *Acting as Secondary* state is converted to a Secondary of the new Primary. The new Primary now starts replaying the DCM to update the Secondary with the writes that were written to the DCM.
- 2 If you do not want to use the **Enable Fast-Failback Logging** option, clear the checkbox, and click **OK** to perform Take Over without the fast-failback logging.

After takeover is complete, to add the Secondary hosts of the original Primary as Secondary hosts of the new Primary, delete the existing RVGs of the original Secondary hosts and then add them as a part of the new Primary.

- 3 If you have chosen to perform the Take Over operation without using fast-failback logging and the original Primary becomes available again, convert it to a Secondary using the **Make Secondary** option. Then resynchronize the original Primary with the new Primary using the **Synchronize Automatically** option. Depending on the size of the data volume this may take quite a while.
Only after the synchronization is complete can you migrate the Primary role back to the original Primary.
After takeover, the existing Secondary becomes the new Primary.
- 4 Assign drive letters to the volumes on the new active site. Ensure that these drive letters are the same as those of the original Primary Site.
- 5 Bring the EVS group online.

Now you can start using Exchange on the new active site (which was the Secondary Site).

Restoring the primary site

After a disaster if the original Primary Site becomes available again you may want to revert the role of the Primary back to this host. Perform the following tasks.

To restore the Primary Site

- 1 Take the RVG resource offline on both the clusters.
- 2 Depending on whether you performed **Takeover** with or without fast-failback option do one of the following:
 - For **Takeover** with the fast-failback option:
The original Primary Site, after it has recovered will be in the `Acting as Secondary` state. If the original system on the Primary Site is not in the `Acting as Secondary` state, verify whether your network connection has been restored.
To synchronize this original Primary and the new active site use the **Resynchronize Secondaries** option from new Primary Site system's right-click menu.
 - For **Takeover** without the fast-failback option:
After performing a takeover without fast-failback you must convert the original Primary to a Secondary using the **Make Secondary** option. Before performing the **Make Secondary** operation the original Primary's RVG and the new Primary's RVG will be shown in separate RDSs. However, after this operation they will be merged under a single RDS.

After the **Make Secondary** operation the original Primary will be converted to a secondary. Right-click on this secondary RVG and select **Start Replication** with **Synchronize Automatically** option.

- 3 After the sychronization is complete, perform a migrate operation to transfer the Primary role back to the original Primary. To do this, right-click on the Primary RVG and select **Migrate** option from the menu that appears.
- 4 Ensure that the volumes have retained the same drive letters that existed before the disaster.
- 5 Bring the RVG resource online on the Secondary.
- 6 Bring the EVS group online on the original Primary.

Deploying SFW and VVR with Microsoft failover clustering: New Exchange installation

This chapter includes the following topics:

- [Tasks for a new Exchange, SFW, VVR, and failover clustering installation \(Windows Server 2008\)](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the primary site](#)
- [Installing SFW with MSCS/Failover Cluster option](#)
- [Setting up the secondary site](#)
- [Installing Exchange on the secondary site](#)
- [Setting up the Exchange group on the secondary site](#)
- [Moving the Volume Manager Disk Group resources to the Exchange group](#)
- [VVR components overview](#)
- [Creating resources for VVR \(primary and secondary sites\)](#)
- [Setting up the replicated data sets \(RDS\) for VVR](#)
- [Creating the RVG resource \(primary and secondary sites\)](#)
- [Setting the database resource dependency on the RVG resource](#)

- [Normal operations and recovery procedures](#)

Tasks for a new Exchange, SFW, VVR, and failover clustering installation (Windows Server 2008)

You can install and configure Storage Foundation for Windows (SFW) and Veritas Volume Replicator (VVR) with Microsoft failover clustering and Exchange 2007 on Windows Server 2008.

After setting up a SFW environment with Microsoft failover clustering for Exchange on a primary site, you can create a secondary or “failover” site using VVR for replication. The example environment describes an active/passive configuration with one to one failover capability.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details about VVR.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 16-1 Tasks for Microsoft failover clustering for Exchange with VVR

Objective	Tasks
“Reviewing the requirements” on page 412	<ul style="list-style-type: none"> ■ Verifying hardware and software prerequisites
“Reviewing the configuration” on page 416	<ul style="list-style-type: none"> ■ Understanding a typical Active/Passive Exchange configuration in a two-node cluster ■ Reviewing the benefits of a dynamic mirrored quorum
Part 1	<ul style="list-style-type: none"> ■ “Configuring the primary site” on page 419.
“Configuring the storage hardware and network” on page 238	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which Exchange will be installed
“Establishing a Microsoft failover cluster” on page 239	<ul style="list-style-type: none"> ■ Reviewing general guidelines to establish a Microsoft failover cluster
“Installing SFW with MSCS/Failover Cluster option” on page 419	<ul style="list-style-type: none"> ■ Installing SFW (automatic installation) ■ Installing the Cluster Option for Microsoft Cluster Service (manual option) ■ Installing the option for Veritas Volume Replicator ■ Configuring VxSAS

Table 16-1 Tasks for Microsoft failover clustering for Exchange with VVR

Objective	Tasks
“Configuring SFW disk groups and volumes” on page 248	<ul style="list-style-type: none"> ■ Using the VEA console to create disk groups ■ Using the VEA console to create the data and log volumes ■ Mounting and unmounting volumes
“Implementing a dynamic mirrored quorum resource” on page 256	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group for the quorum resource with a mirrored volume ■ Adding the Volume Manager Disk Group resource for the quorum ■ Configuring the cluster quorum settings and changing the quorum resource to a dynamic mirrored quorum resource
“Adding a Volume Manager Disk Group resource for Exchange 2007 installation” on page 259	<ul style="list-style-type: none"> ■ Adding a Volume Manager Disk Group resource for the SFW disk group that was created for the First Storage Group. You add it to the quorum group and then later move it to the Exchange group. This allows the First Storage Group to be installed to a dynamic disk.
“Installing Exchange Server” on page 259	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Running the Microsoft Exchange Server installation ■ Repeating the procedure on additional nodes
“Adding the Volume Manager Disk Group resources to the Exchange group” on page 260	<ul style="list-style-type: none"> ■ Moving the Volume Manager Disk Group resource for the First Storage Group to the Exchange Group ■ Adding any other Volume Manager Disk Group resources needed for the Exchange databases
“Moving Exchange databases and logs to shared storage” on page 261	<ul style="list-style-type: none"> ■ If necessary, altering the paths for the databases and transaction log
“Verifying the cluster configuration” on page 263	<ul style="list-style-type: none"> ■ Moving the online cluster group to the second node and back to the first node
Part 2	“Setting up the secondary site” on page 430. Creating a parallel configuration at the secondary site, with some differences in procedure for installing Exchange and setting up the Exchange group.
Part 3	Adding the components for VVR replication. See “VVR components overview” on page 434.

Table 16-1 Tasks for Microsoft failover clustering for Exchange with VVR

Objective	Tasks
“Creating resources for VVR (primary and secondary sites)” on page 434	<ul style="list-style-type: none"> ■ Using Failover Cluster Management to create the Network Name and IP address for an RVG ■ Creating the Replicator Log volumes for VVR
“Setting up the replicated data sets (RDS) for VVR” on page 435	<ul style="list-style-type: none"> ■ Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites
“Creating the RVG resource (primary and secondary sites)” on page 446	<ul style="list-style-type: none"> ■ Using Failover Cluster Management to set up an RVG resource
Setting the database resource dependency on the RVG resource	<ul style="list-style-type: none"> ■ Using Failover Cluster Management to set up the database resource dependency on the RVG resource
“Normal operations and recovery procedures” on page 449	<ul style="list-style-type: none"> ■ Operating the Microsoft failover cluster with VVR and recovering from a disaster

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

Supported software for Microsoft failover clusters with SFW

Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

The following software is supported for deploying Microsoft Exchange with SFW and Microsoft clustering on Windows Server 2008:

- Veritas Storage Foundation 5.1 Service Pack 2 for Windows (SFW)
 Include the following option during installation:
 - Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster
 For a DR configuration with Veritas Volume Replicator, include the following option:
 - Veritas Volume Replicator option

The following table lists the Microsoft Exchange Server versions supported on Windows Server 2008 with this release of SFW for a Microsoft clustering solution.

Note: For Exchange 2007, you can only cluster the Mailbox server role. Refer to the Microsoft documentation for other Exchange requirements.

Table 16-2 Supported Microsoft Exchange Server versions

Exchange Server	Windows Servers
Microsoft Exchange Server 2003 Standard Edition or Enterprise Edition on Windows Server 2003 (Exchange Server 2003 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 (32-bit) R2 Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft Exchange Server 2007 (SP1, SP2, or SP3), Standard Edition or Enterprise Edition on Windows 2008 Server (Mailbox server role required)	<ul style="list-style-type: none"> ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition ■ Windows Server 2008 x64 R2 without Hyper-V on Standard, Enterprise, Datacenter Editions ■ Windows Server 2008 x64 R2 on Standard, Enterprise, Datacenter Editions (for physical host or guest, not parent partition/Hyper-V integration) ■ Windows Server 2008 R2 for IA Systems - IA64 ■ Windows Server 2008 x64 R2 Web Edition ■ Windows Server 2008 on all current editions and architectures Symantec currently supports (SP2 required) ■ Windows Storage Server 2008

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space.

The following table summarizes disk space requirements for SFW.

Table 16-3 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

System requirements

Observe the following system requirements:

- One CD-ROM drive accessible to the system on which you are installing SFW.
- The configuration described requires shared disks to support applications that migrate between nodes in each cluster.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- Microsoft failover clustering requires at least two network adapters per system (one NIC to connect each system to the public network, and one NIC for the private network on each system). Symantec recommends using three network adapters (two NICs exclusively for the private network and one for the public network). Route each private NIC through a separate hub or switch to avoid single points of failure.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. There are six network interface cards, three for each server (two each for the private network and one for the public network). You also need a static IP address for the cluster itself.
- Each system requires a minimum of 1 GB of RAM for SFW.
- A minimum of 2 GB of RAM per server is required for Exchange 2007; refer to your Microsoft documentation for more information.
- Systems to be clustered must be configured as part of a Windows Server 2008 domain. Each system in a Microsoft failover cluster must be in the same domain and must be using the same operating system version.
- To install and work with the SFW and Microsoft clustering software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- The administrator account for the Exchange virtual server (clustered mailbox server) on the primary site must be the same account used for the Exchange virtual server on the secondary site.
- Verify that the DNS and Active Directory Services are available. Make sure a reverse lookup zone exists in the DNS. Refer to the Microsoft Exchange documentation for instructions on creating a reverse lookup zone.
- Symantec recommends two disks for Exchange: one for Exchange database files and one for Exchange log files.
- The cluster on the secondary site must reside in the Active Directory domain of the cluster on the primary site.

Note: Refer to the Hardware Compliance List on the Symantec Support web site at <http://www.symantec.com/business/support/index.jsp> to determine the approved hardware for SFW.

Reviewing the configuration

The example configuration is for high availability within a cluster, and a clustered configuration with VVR replication between two sites. In an active/passive configuration, one or more Exchange clustered mailbox servers (previously called virtual servers) can exist in a cluster. In this case, Exchange clustered mailbox server EVS1 can fail over from SYSTEM1 to SYSTEM2 on the primary site, and SYSTEM5 to SYSTEM6 on the secondary site.

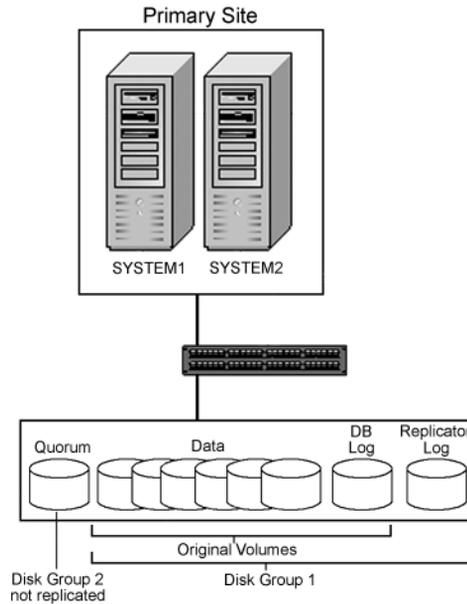
Each site has a cluster with the nodes set up appropriately for failover within the site.

The cluster on the primary site has an SFW cluster disk group that contains the volumes required by VVR for setting up a Replicated Volume Group (RVG). The Exchange application data is stored on the volumes that are under the control of the RVG.

You create a separate disk group on each site for the cluster quorum volume, which is not replicated, because each cluster has its own quorum.

[Figure 16-1](#) shows a cluster configuration on the primary site.

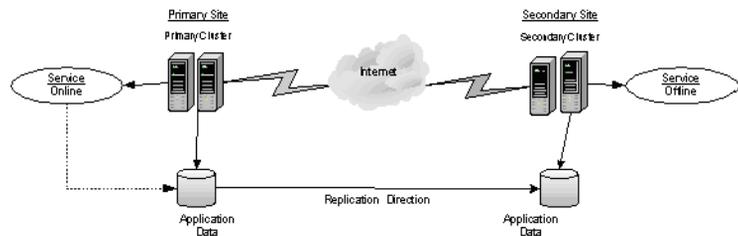
Figure 16-1 DR configuration on the primary site



In a two-site VVR environment, the cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the primary cluster fails. The following illustration displays an environment that is prepared for a disaster with a two-site VVR solution. In this case, the primary site is replicating its application data to the secondary site.

Figure 16-2 shows an example disaster recovery configuration before a failure.

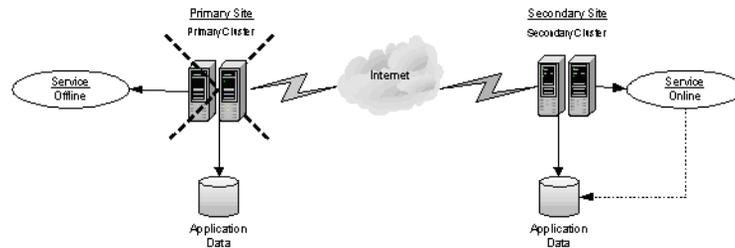
Figure 16-2 Disaster recovery example configuration



When a failure occurs (for instance, after an earthquake that destroys the data center in which the primary site resides), the data that was replicated to the secondary site can be used to restore the application services to clients.

Figure 16-3 shows an example disaster recovery configuration after a failure.

Figure 16-3 Disaster recovery configuration after a failure



The following are some other key points about the configuration:

- A Microsoft failover cluster must be running before you can install SFW.

Note: You can install the SFW option for MSCS/Failover Cluster on a machine that is not a member of a Microsoft cluster. However, if that machine becomes the first node in a Microsoft cluster, the Volume Manager Disk Group resource type must be manually registered. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

- SFW enables you to create a dynamic mirrored quorum. The quorum resource maintains the cluster database and critical recovery information in a recovery log.

Microsoft clustering only supports a basic physical disk and does not enable you to mirror the quorum resource. One advantage of SFW is that it provides a dynamic mirrored quorum resource for Microsoft clustering. If a quorum disk fails, a mirror on another disk (another plex) takes over and the resource remains online. In this configuration, Symantec recommends creating a three-way mirror for the quorum to provide additional fault tolerance. If possible, do not use the disks assigned to the quorum for any other purpose.

You can wait until the end of setting up the environment to convert the basic physical disk quorum into a dynamic mirrored volume; this enables you to verify that Exchange is working in the cluster before adding the dynamic quorum volume.

- VVR does not support SFW (software) RAID 5 volumes.

Configuring the primary site

The beginning steps for configuring the primary site for a disaster recovery solution are the same as when configuring for high availability.

You begin with the following topics in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):

- [“Configuring the storage hardware and network”](#) on page 238
- [“Establishing a Microsoft failover cluster”](#) on page 239

You can then continue with the following procedure for installing SFW for a disaster recovery configuration. In configuring the primary site to prepare for a disaster recovery rather than for high availability alone, you need to make sure that the VVR option is selected during the SFW installation. After installation you must configure the Veritas Volume Replicator Security Service (VxSAS).

Installing SFW with MSCS/Failover Cluster option

This section assumes that you are running a Microsoft failover cluster and that you are installing SFW on an inactive system that does not own any cluster resources.

Symantec recommends a rolling installation to install SFW. For a rolling installation, you must first install SFW on an inactive system. After SFW is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then install SFW on the other inactive systems in the Microsoft failover cluster simultaneously.

SFW installation tasks

The product installer enables you to install the software for Veritas Storage Foundation 5.1 for Windows. The installer automatically installs SFW. You must select the options to install VVR, and the Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation. The steps in this section are based on a server installation.

Installing SFW involves the following:

- Performing pre-installation tasks
See [“Pre-installation tasks”](#) on page 420.
- Installing the product

See “[Installing Veritas Storage Foundation for Windows](#)” on page 420.

- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 424.

Pre-installation tasks

Perform the following pre-installation task:

- Moving the Online Groups
See “[Moving the online groups](#)” on page 420.

Moving the online groups

If your resource groups are online on the system where you are installing SFW. You must move the resource groups from the SFW system to another system in the cluster.

To move the online groups

- 1 Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click a resource group and then click **Move this service or application to another node > Move to node [name of node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were moved. This confirms that the resource groups have moved to another system.
- 4 If you need to move the resource groups back to the original system, repeat Step 2.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install Veritas Storage Foundation for Windows on a MSCS configuration.

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 3 Review the links on the DVD browser panel.
The panel provides the **Late Breaking News** link to access the latest information about updates, patches, and software issues regarding this release, and a link to run the Configuration Checker to verify that your configurations meet all pertinent software and hardware requirements. The panel provides links to install the software (Storage Foundation for Windows or Storage Foundation HA for Windows) and access the documentation (Getting Started Guide, Installation and Upgrade Guide, and Release Notes).
The panel also provides links to access the Veritas Operations Services (VOS) site (VOS provides you four types of detailed reports about your computer and Symantec enterprise products, a checklist of configuration recommendations, and system and patch requirements to install or upgrade your software), contact the Symantec Technical Support, and see the contents of the DVD.
- 4 Under Install Storage Foundation, do one of the following:
 - Click the **Complete/Custom** link to install server or client or both the components.
 - Click the **Administrative Console** link to install only the client components.

Note: With the Administrative Console option, you will not be prompted for a product license or presented with a list of product options for SFW or SFW HA.

Click the **Complete/Custom** link.

- 5 On the Welcome panel, review the Welcome message and the listed prerequisites. Ensure that the prerequisites are met prior to proceeding. Click **Next**.
- 6 On the License Agreement panel, read the license agreement. If you agree to the license terms, click **I accept the terms of the License Agreement**, and then click **Next**.
- 7 On the License panel, enter the product license key before adding license keys for features. Click **Enter license key(s)**, provide the license key in the field below it, and then click **Add**.
If you do not have a license key, click **Use embedded evaluation license key** to use the default evaluation license key. This license key is valid only for a limited evaluation period.
To remove a license key, click the key, and then click **Remove**. To see a license key's details, click the key to display its details in the License key details area.
Click **Next** to continue.
- 8 On the Option Selection panel, do the following, and then click **Next**:
 - Select the **Volume Replicator (VVR)** option.
 - Select the **Cluster Option for Microsoft Cluster Service (MSCS)/Failover Cluster** option.
 - Select any additional options applicable to your environment.
 - Make sure that the Client Components option is selected to install the client component.
- 9 On the System Selection panel, do the following, and then click **Next**:
 - To add a computer for installation, provide the name of the computer in the System Name box.
OR
If you do not know the name of the computer, click **Browse** to search for the computers available in your domain. The Select Systems dialog box appears. Select a computer from the Available Systems area, move it to the Selected Systems area, and then click **OK** to add it for installation.
 - To change the installation path of an added computer, click the folder icon for the computer, and then select the installation path in the Browse For Folder dialog box.
 - To know the verification status and other information of the added computer, click the information icon.
 - To remove an added computer, select it, and then click the recycle bin icon.

Note: When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.

The installer checks the prerequisites for the added computers and displays the results in the Status column. If a computer fails validation, address the issue, and repeat the validation process by clicking **Re-verify**.

- 10 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages. If applicable to your installation, perform the procedure mentioned in the messages. Review the messages, and then click **OK**.

- **Quorum Arbitration**

The quorum arbitration settings are used to set the time that Microsoft clustering allows for quorum arbitration. You may want to adjust the minimum and maximum time for quorum arbitration to ensure optimal functioning of Veritas Storage Foundation for Windows dynamic volumes with MSCS. 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. For more information, see the *Veritas Storage Foundation for Windows Administrator's Guide*.

- **Dynamic Multi-pathing**

If you are using multiple paths and select a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical connection during installation. After the installation, reconnect additional physical paths before rebooting the system.

- 11 On the Pre-install Summary panel, the Pre-install Report is displayed with summarized information about the installation. Review the Pre-install Report. Click **Back** to make changes, if necessary. Click **Save Report** to save the report as a web page or text file on your computer.

It is recommended that you select the **Automatically reboot systems after installer completes the operation** check box to restart the computer after the installation is complete.

Click **Install** to install the software.

- 12 The Installation panel displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report, and address the reason for failure. You may have to either repair the installation or uninstall and re-install the software.
- 13 On the Post-install Summary panel, the Post-install Report is displayed with summarized information about the installation results along with links to the log files and installation summary for the computer. Click **Save Report** to save the report as a web page or text file on your computer. Review the Post-install Report and log files, and then click **Next**.
- 14 On the Finish panel, click **Finish** to complete the installation.
- 15 Click **Yes** to restart the local node.

Post-installation tasks

You must perform the following post-installation tasks:

- Moving the Online Groups
See “[Moving the online groups](#)” on page 424.
- Completing the SFW Installation for the other systems in the failover cluster
See “[Completing the SFW installation](#)” on page 425.
- Configuring the VxSAS service
See “[Configuring the VxSAS service](#)” on page 425.

Moving the online groups

You can move the resource groups from the system, back to the previous system where SFW is installed.

To move the online groups

- 1 Open the Failover Cluster Management tool. (**Start > Administrative Tools > Failover Cluster Management**).
- 2 In the left pane, under Services and Applications, right-click the resource group and then click **Move this service or application to another node > Move to node [name of original node]**.
If there is more than one resource group, you must repeat this step until all the resource groups are moved back to the original node.
- 3 In the Failover Cluster Management console, center panel, verify that the Current Owner name has changed for all of the resource groups that were

moved. This confirms that all the resource groups have moved back to the original system.

Completing the SFW installation

You must repeat the SFW installation tasks for the other systems in the MSCS cluster.

See “[SFW installation tasks](#)” on page 419.

Configuring the VxSAS service

Complete the following procedure to configure this required service for VVR. This procedure should not be done until you have installed SFW on all cluster systems. Otherwise, you will get an error message from the VxSAS wizard if you try to select a system without SFW installed.

You can run the VxSAS wizard from any site once SFW is installed on all cluster systems; at that time, you can run the wizard for both the primary and secondary site systems. The MSCS groups can be either online or offline.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

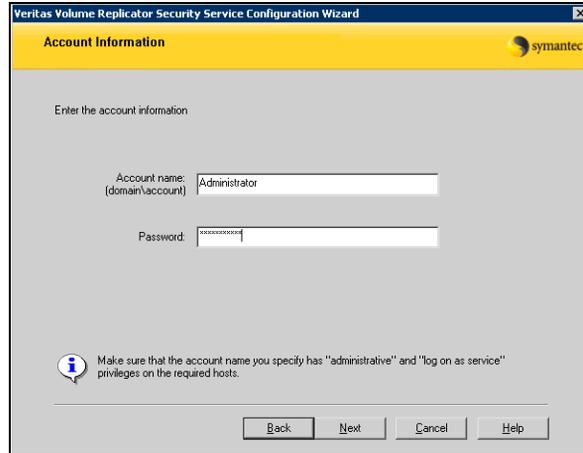
The welcome page appears. This page displays important information that is useful as you configure the VxSAS service. Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information wizard page as follows:

Account name (domain\account) Enter the administrative account name in the Account name field.

Password Specify a password in the **Password** field.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, then make sure you specify the same user name and password when configuring the VxSAS service on the other hosts.



- 3 After providing the required information, click **Next**.
- 4 Select the required domain to which the hosts that you want to configure belong, from the Domain Selection wizard page.

Selecting Domains The Available Domains pane lists all the domains that are present in the Windows network neighborhood.
Select the required domain by moving the appropriate name from the Available Domains pane to the Selected Domains pane, either by double-clicking it or using the arrow button.

Adding a Domain If the domain name that you require is not displayed, then add it by using the **Add Domain** option. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected Domains list.

After specifying the domain, click **Next**.

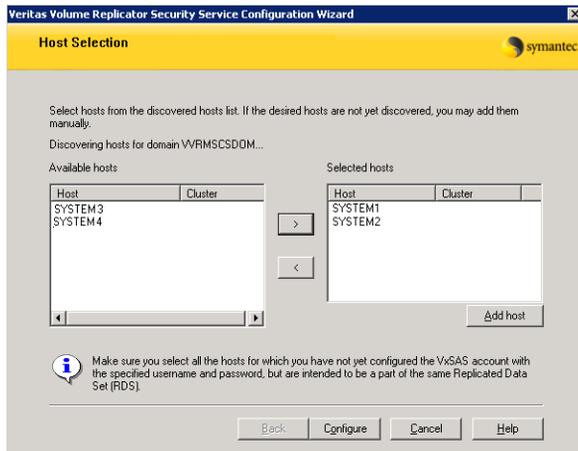
5 Select the required hosts from the Host Selection page.

Selecting Hosts The Available Hosts pane lists the hosts that are present in the specified domain.

Select the required host by moving the appropriate name from the Available Hosts list to the Selected Hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a Host If the host name you require is not displayed, then add it using **Add Host** option. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected Hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.



- 6 After the configuration completes, the Configuration Results page is displayed. If the operation is successful, then the Status column displays the appropriate message to indicate that the operation was successful.

If the operation was not successful, then the page displays the status as failed and the corresponding details on why the account update failed. It also displays the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

- 7 Click **Finish** to exit the wizard.

After completing SFW installation on the primary site, continue with the additional tasks for completing the primary site configuration.

See “[Completing the primary site configuration](#)” on page 429.

Completing the primary site configuration

After you have installed SFW, the remaining steps for configuring the primary site for disaster recovery are the same as when configuring high availability.

To complete the primary site configuration, refer to the following topics in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):

- [“Configuring SFW disk groups and volumes”](#) on page 248
- [“Implementing a dynamic mirrored quorum resource”](#) on page 256
- [“Adding a Volume Manager Disk Group resource for Exchange 2007 installation”](#) on page 259
- [“Installing Exchange Server”](#) on page 259
- [“Adding the Volume Manager Disk Group resources to the Exchange group”](#) on page 260
- [“Moving Exchange databases and logs to shared storage”](#) on page 261
- [“Verifying the cluster configuration”](#) on page 263

Setting up the secondary site

After setting up an SFW environment with Microsoft failover clustering and Exchange 2007 on the primary site, configure the secondary site. Many tasks are the same as on the primary site.

The following is the list of tasks in the recommended order:

- Configure the storage hardware and network and establish a Microsoft failover cluster.
Refer to the following topics in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):
 - [“Configuring the storage hardware and network”](#) on page 238
 - [“Establishing a Microsoft failover cluster”](#) on page 239
- Install SFW for a disaster recovery configuration and configure the VxSAS service. Refer to the following topic in this chapter:
 - [“Installing SFW with MSCS/Failover Cluster option”](#) on page 419
- Configure SFW disk groups and volumes.
During the creation of disk groups and volumes for the secondary site, make sure the following are exactly the same as the cluster on the primary site:
 - Cluster disk group names
 - Volume names and sizes
 - Drive lettersRefer to the following topic in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):
 - [“Configuring SFW disk groups and volumes”](#) on page 248
- Implement a dynamic mirrored quorum resource
Refer to the following topic in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):
 - [“Implementing a dynamic mirrored quorum resource”](#) on page 256
- Add Volume Manager Disk Group resources for the Exchange 2007 disk group(s)
The resource will not be used during the Exchange 2007 installation but will be used in the process of recovering the existing database and log information from Active Directory after Exchange installation.
If you created any other Exchange storage groups, include resources for those disk groups as well.
Refer to the following topic in [Chapter 12, “Deploying SFW with Microsoft failover clustering: New Exchange installation”](#):

- [“Adding a Volume Manager Disk Group resource for Exchange 2007 installation”](#) on page 259
- Install Exchange 2007
The procedure for installing and configuring Exchange 2007 on the secondary site is different from on the primary site. See the following topic in this chapter:
 - [“Installing Exchange on the secondary site”](#) on page 431
- Set up the Exchange group on the secondary site
Unlike on the primary site, Exchange installation does not create the Exchange group and resources. However, after Exchange installation, you can run a command that recovers the resource information from Active Directory. See the following topic in this chapter:
 - [“Setting up the Exchange group on the secondary site”](#) on page 432
- Move the Volume Manager Disk Group resources to the Exchange group
Move the resources you added for Exchange to the quorum group to the Exchange group. See the following topic in this chapter:
 - [“Moving the Volume Manager Disk Group resources to the Exchange group”](#) on page 433

After completing these tasks you will have a clustered site with SFW with the Microsoft cluster option and Exchange installed on all the nodes on both the primary and the secondary sites.

Installing Exchange on the secondary site

Exchange 2007 requires service pack 1 for Windows Server 2008.

Review the following requirements for installing Exchange Server 2007 on the secondary site:

- Before you begin installation, make sure to take the Exchange cluster group offline on the primary site; otherwise, the installation on the secondary site will not function properly.
- On the secondary site nodes, do not install the Active Clustered Mailbox role. Instead, install the Passive Clustered Mailbox role on all nodes of the secondary site.
- Make sure that the same drive letter and path are available on all nodes and have adequate space for the installation. For example, if you install the Exchange program files in `C:\Program Files\ExchSrvr` on one node, you must install the files in `C:\Program Files\ExchSrvr` on all the other nodes. Use a drive letter on the local disk of each node in the cluster.

Refer to the Microsoft documentation for additional prerequisites and details on Exchange installation.

Setting up the Exchange group on the secondary site

Installing Exchange 2007 on the secondary site does not create the Exchange group and resources as it does on the primary site. However, after Exchange installation, you can use run the RecoverCMS command to recover the Exchange group and resource information from Active Directory.

To set up the Exchange group on the secondary site

- 1 On the secondary site, ensure that you have set up the cluster disk groups and dynamic volumes to match those on the primary site.
- 2 Ensure that you have created the Volume Manager Disk Group resources that correspond with the Exchange disk groups.
- 3 Ensure that the disk groups are online and volumes mounted on the node where you are going to run the RecoverCMS command.
- 4 Change to the directory where you installed Exchange.
- 5 Run the RecoverCMS command, using the following syntax:

```
Setup.com /recoverCMS  
/CMSName:<name>/CMSIPAddress:<ip>
```

Where *<name>* is the name you assigned to the Exchange server during installation on the primary site and *<ip>* is the IP assigned during Exchange installation on the primary site.

You may receive an error message due to a Microsoft issue that states the command failed to bring cluster resource Network Name online, and that the group or resource is not in the correct state to perform the requested operation.

In this case, complete the following steps on the secondary site:

- On the secondary site, using the Failover Management console, bring the clustered mailbox server (Network Name) resource online manually.
- After the Network Name resource comes online, take the IP resource offline. (This will also offline the Network Name resource.)
- Delete the Network Name resource and the IP resource.
- Run the RecoverCMS command over again. This time it will succeed and all the Exchange resources will be created on the secondary site.

- 6 On the secondary site, bring the Exchange cluster group and its resources offline.
- 7 On the primary site, bring the Exchange cluster group and its resources online.

Moving the Volume Manager Disk Group resources to the Exchange group

You originally added the Volume Manager Disk Group (VMDG) resources for Exchange to the cluster quorum group so that they could be used in setting up the Exchange group on the secondary site. You now move those resources to the Exchange group.

To move the Volume Manager Disk Group resources to the Exchange group

- 1 In the quorum group, right-click the resource you created for Exchange and select the option to move it to another group.
- 2 Select the Exchange group as the target for the move and click **OK**.
- 3 Repeat these steps for any other resources you created for Exchange.

VVR components overview

The next set of tasks is to configure both sites for VVR replication. You configure the following Veritas Volume Replicator components:

Replicated Volume Group (RVG)	<p>An RVG is made up of one or more volumes in a SFW disk group. The updates made on the RVG on the primary host are sent to a configured secondary host. Thus, on the secondary host 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. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG.</p> <p>An RVG within a disk group is the container for replication, so if you have multiple disk groups, you will 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.</p>
Replicated Data Set (RDS)	<p>An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).</p>
Replicator Log volume	<p>Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any RVG updates 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 needs to become the new primary site. The log volumes at the two sites must have the same name. Symantec recommends having Replicator Log volumes of the same size at the primary site and the secondary site.</p>

Creating resources for VVR (primary and secondary sites)

Create the resources for VVR replication at the primary and secondary sites using the Failover Cluster Management tool. You create a network name resource and IP address resource to be used for VVR replication.

A separate valid IP address is necessary for VVR replication, because on the secondary cluster before a disaster, the application IP must be offline whereas the VVR IP must be online.

You create the resources for the primary site and then repeat the procedure to create the resources on the secondary site.

To create a Network Name resource and IP address resource for VVR replication

- 1 Right-click on the application group and select **Add a Resource > Client Access Point**.
- 2 In the Client Access Point panel of the New Resource Wizard, specify the following:
 - In the **Name** field, specify a name for the Network Name resource. The default is the name of the group you selected. Specify any name except the node and the virtual server name. The network name you assign when creating the resource for the secondary site must be different from the network name for the primary site.
 - Select the network and specify the IP address.Click **Next**.
- 3 In the Confirmation panel, review the information and click **Next**.
- 4 When configuration is complete, click **Finish**.
- 5 Repeat the same procedure to create the IP and the Network Name resource at the secondary site.
- 6 Bring the resources online.

Setting up the replicated data sets (RDS) for VVR

For each disk group you created for the application, you set up a Replicated Data Set (RDS) on the primary and secondary hosts. The Setup Replicated Data Set Wizard enables you to configure an RDS for both sites.

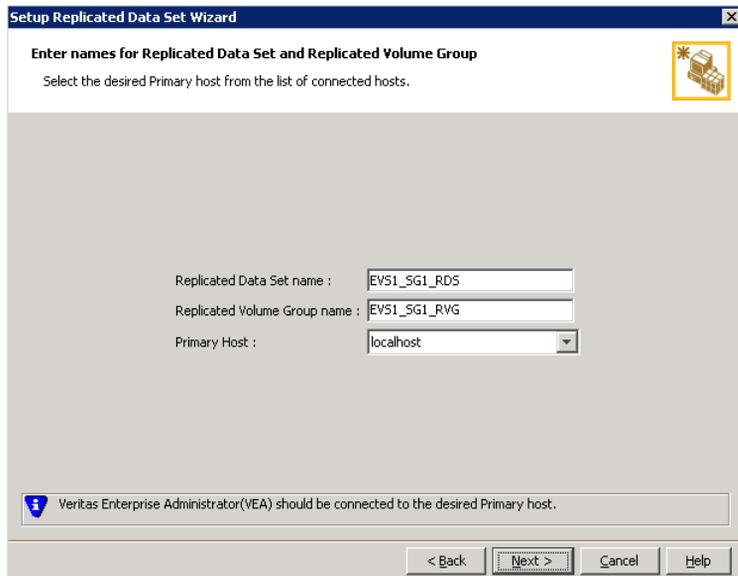
- Verify that the set of volumes that are intended to store the user Exchange database files, the user Exchange log files, and the Replicator log file have been created on the Primary node.
- Verify that the EVS IP resource is offline on the Secondary site. This would also offline all the dependent Exchange resources.
- Verify that the data and replicator log volumes are not of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volume names containing a comma
- Verify that the Replicator Log volume does not have a DCM.

- Verify that the cluster disk group is imported on the primary and secondary site.
- Verify that VxSAS has been configured.

The following procedure enables you to set up an RDS on the primary and secondary sites and start replication.

To create the Replicated Data Set

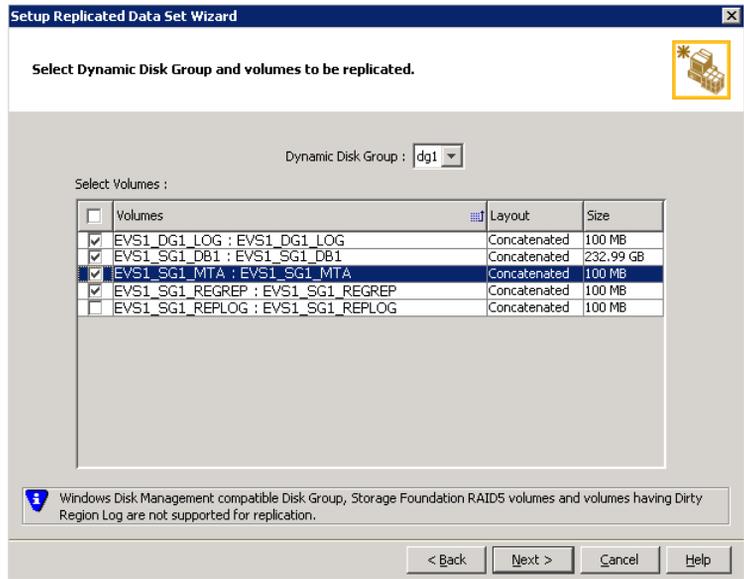
- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Setup Replicated Data Set**.
- 3 Read the information on the Welcome page and then click **Next**.
- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG) and then click **Next**.



By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

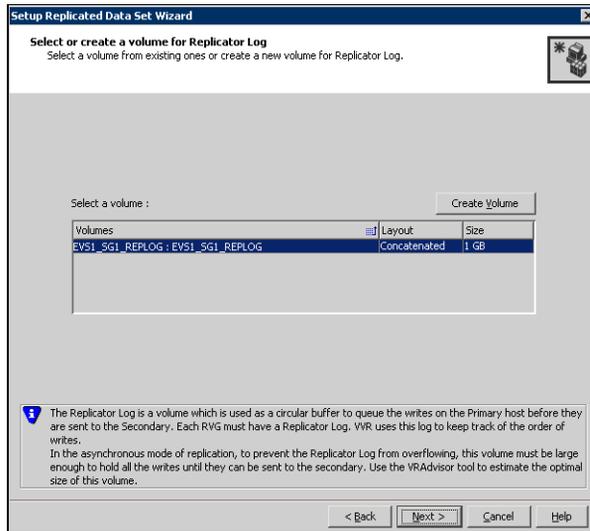
- 5 Select from the table the dynamic disk group and data volumes that will undergo replication and then click **Next**.



To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

6 Complete the select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (EVSI_SG1_REPLOG).
If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

- Name** Enter the name for the volume in the **Name** field.
- Size** Enter a size for the volume in the **Size** field.
- Layout** Select the desired volume layout.

Disk Selection

Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** check box to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The check box will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this check box along with **Select disks manually** option, then the user can select only TP disks from **Available Disks**.

For more information on Thin Provisioning refer to the *Veritas Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want VVR to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on it or select **Add** to move the disks into the Selected disks pane.

- Click **OK** to create the Replicator Log volume.
- Click **Next** in the **Select or create a volume for Replicator Log** dialog box.

7 Review the information on the summary page and click **Create Primary RVG**.

8 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

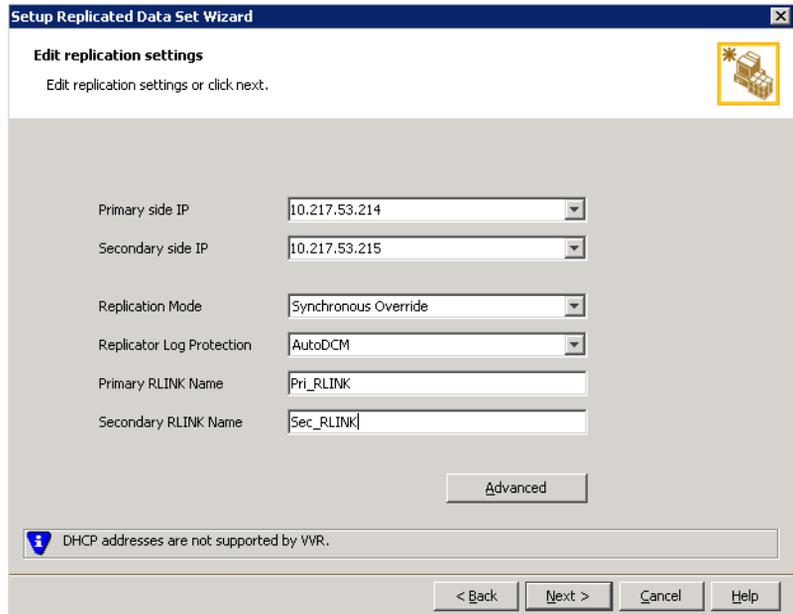
9 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field and then click **Next**.

If the Secondary host is not connected to VEA, the wizard tries to connect it when you click Next. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the Add Secondary option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 10 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.
The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:
 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary
Otherwise, the RDS setup wizard enables you to create the required volumes manually.
 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
 - 11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.
This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.
 - If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
 - If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.
Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.
When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.
- If all the data volumes to be replicated meet the requirements, this screen does not occur.

12 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:



- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Secondary side IP Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode	<p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as <code>MISSING</code>.</p>
Replicator Log Protection	<p>The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.</p> <p>The DCM option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.</p> <p>The Off option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to Off, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.</p>

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK
Name

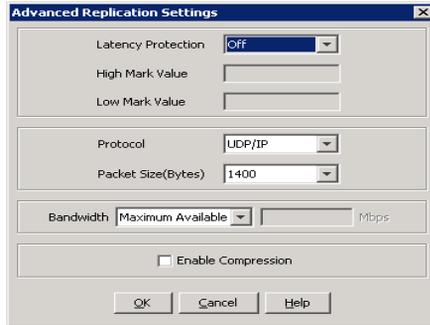
This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Secondary RLINK
Name

This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

- Click **Next** to start replication with the default settings.

- 13 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

Off is the default option and disables latency protection.

Fail enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.

Override enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Enable Compression Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box and then click **Next**.

14 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from Checkpoint If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu.
- Click **Next** to display the Summary page.

15 Review the information.

Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the RVG resource (primary and secondary sites)

To enable a disaster recovery setup, once VVR is configured you will need to create an RVG resource on the primary and secondary sites.

To create a Replicated Volume Group (RVG) resource

- 1 In Failover Cluster Management, expand Services and Applications, right-click the Exchange Virtual Server (EVS) group that you have created and select **Add a resource > More resources > Add Replicated Volume Group**.
The New Replicated Volume Group appears in the center panel under Disk Drives.
- 2 Right-click **New Replicated Volume Group** and click **Properties**.
- 3 On the General tab of the Properties dialog box, in the Resource Name field, type a name for the RVG resource.
- 4 On the Dependencies tab, add the dependencies for the RVG resource:
 - Click the box **Click here to add a dependency**

- From the Resource drop-down list, select the network name you created for the RVG. Click **Insert**.
 - Click the box **Click here to add a dependency**
 - From the Resource drop-down list, select the Volume Manager Disk Group resource created for the application disk group. Click **Insert**.
- 5 On the Properties tab, specify the following:
 - In the rvgName field, type the same name that you assigned the RVG on the General tab.
 - In the dgName field, type the name assigned in the VEA to the application disk group.
 - 6 Click **OK** to close the Properties dialog box.
 - 7 Right-click the RVG resource and click **Bring this resource online**.
 - 8 Repeat the same steps to create the RVG resource at the secondary site.

Setting the database resource dependency on the RVG resource

The Exchange database resource was earlier set to depend on a Volume Manager Disk Group resource that corresponded to the disk group created for the application. After you add the RVG resource for that disk group, you must change the dependency. You set the database resource to depend on the RVG resource instead.

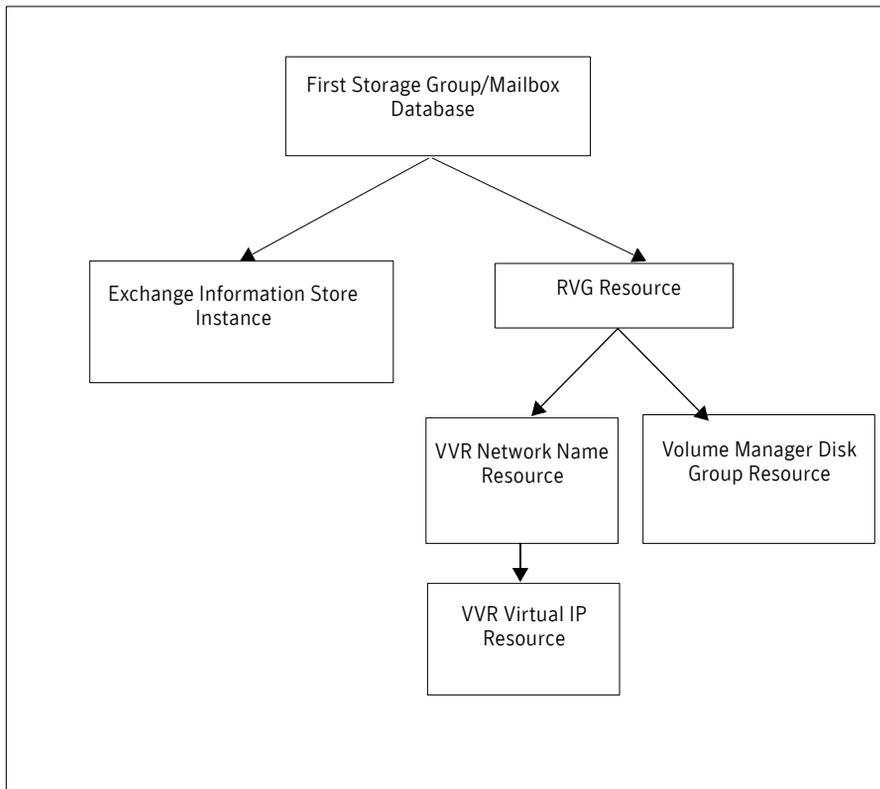
To set the database resource dependency on the RVG resource

- 1 In Failover Cluster Management, select the Exchange resource group.
- 2 In the result pane, under Other Resources, right-click the appropriate database resource and select Properties.
- 3 In the Dependencies tab of the Properties dialog box:
 - Click the box **Click here to add a dependency**.
 - Select the Replicated Volume Group resource from the dropdown list of available resources.
 - Select the Volume Manager Disk Group (VMDG) resource from the dependencies list and click **Delete**.
- 4 Click **OK** to close the Properties dialog box.
- 5 Repeat these steps for any additional database resources.

The cluster configuration is now complete. Bring online the entire application group on the primary cluster.

Figure 16-4 shows a dependency graph with the dependencies that have been established for the RVG resource.

Figure 16-4 Example dependency graph with a replicated volume group



Normal operations and recovery procedures

This section gives considerations for normal VVR operations and also describes the recovery process.

Normal operations

Under normal operating conditions you can monitor the status of the replication using the following:

- VEA GUI
- Command Line Interface (CLI)
- Performance Monitor (perfmon)
- Alerts

For details, refer to the “Monitoring Replication” section in the *Veritas Volume Replicator Administrator’s Guide*.

Performing planned migration

For maintenance purposes, or for testing the readiness of the Secondary host you may want to migrate the application to the Secondary host. These are a generic set of tasks that you may need to perform.

To migrate the application to the Secondary host

- 1 Take the RVG resource offline on both the clusters, which will take all dependent resources (such as the Exchange resources) offline.
- 2 Transfer the Primary role to the Secondary using the Migrate option. From the **VEA** screen, right-click the Primary RVG and select **Migrate**. Select the Secondary host and click **OK**. The replication role is migrated to the Secondary host.
- 3 Assign drive letters to the volumes on the new Primary. Ensure that these drive letters are the same as those of the original Primary.
- 4 Bring the RVG resource online on the new Secondary.
- 5 Bring the EVS group online on the new Primary.

You can now verify that the Exchange runs fine on the new Primary with the replicated data. After verifying you can revert back the roles to its original state using the same set of tasks described above.

Any changes that you make to the data on the new Primary will get replicated to the original Primary, which is now the Secondary.

Replication recovery procedures

This section provides information on bringing up an Exchange server on the Secondary Site, in the event of a disaster. It also explains how to migrate the Primary role back to the original Primary Site once it is in a good state after a disaster.

Bringing up Exchange on the secondary site

To bring up Exchange on the Secondary Site, recover the Exchange data.

To recover the Exchange data

- 1 From the left-pane in the VEA GUI console on a system on the Secondary Site, right-click on the desired secondary RVG node inside the replication network. Select the **Take Over** option. The **Take Over** dialog box is displayed.
 - By default, the **Enable Fast-Failback Logging** option is selected if the data volumes have DCM logs associated with them. You can use this option to perform takeover with fast-failback logging.

The DCM is activated for fast-failback logging and the new incoming writes are marked on the DCM of the New Primary.

If the replication status of Secondary RVG was *Inactive* when the Primary failed, then the **Enable Fast-Failback Logging** option is unavailable for selection. In this case you can perform **Take Over** without using fast-failback logging.
 - Select the **Synchronize Automatically** option if you want the new Primary and the original Primary to get synchronized automatically, after the original Primary recovers.

If you have not selected this option, the original Primary, after it recovers will be in the *Acting as Secondary* state. To synchronize this original Primary with the new Primary use the **Resynchronize Secondaries** option from new Primary RVG's right-click menu. When the resynchronization starts, the original Primary which was in the *Acting as Secondary* state is converted to a Secondary of the new Primary. The new Primary now starts replaying the DCM to update the Secondary with the writes that were written to the DCM.
- 2 If you do not want to use the **Enable Fast-Failback Logging** option, clear the checkbox, and click **OK** to perform Take Over without the fast-failback logging.

After takeover is complete, to add the Secondary hosts of the original Primary as Secondary hosts of the new Primary, delete the existing RVGs of the original Secondary hosts and then add them as a part of the new Primary.

- 3 If you have chosen to perform the Take Over operation without using fast-failback logging and the original Primary becomes available again, convert it to a Secondary using the **Make Secondary** option. Then resynchronize the original Primary with the new Primary using the **Synchronize Automatically** option. Depending on the size of the data volume this may take quite a while.
Only after the synchronization is complete can you migrate the Primary role back to the original Primary.
After takeover, the existing Secondary becomes the new Primary.
- 4 Assign drive letters to the volumes on the new active site. Ensure that these drive letters are the same as those of the original Primary Site.
- 5 Bring the EVS group online.

Now you can start using Exchange on the new Primary, which was the Secondary Site.

Restoring the primary site

After a disaster if the original Primary Site becomes available again you may want to revert the role of the Primary back to this host. Perform the following tasks.

To restore the Primary Site

- 1 Take the RVG resource offline on both the clusters.
- 2 Depending on whether you performed **Takeover** with or without fast-failback option do one of the following:
 - For **Takeover** with the fast-failback option:
The original Primary Site, after it has recovered will be in the `Acting as Secondary` state. If the original system on the Primary Site is not in the `Acting as Secondary` state, verify whether your network connection has been restored.
To synchronize this original Primary and the new active site use the **Resynchronize Secondaries** option from new Primary Site system's right-click menu.
 - For **Takeover** without the fast-failback option:
After performing a takeover without fast-failback you must convert the original Primary to a Secondary using the **Make Secondary** option. Before performing the **Make Secondary** operation the original Primary's RVG and the new Primary's RVG will be shown in separate RDSs. However, after this operation they will be merged under a single RDS.

After the **Make Secondary** operation the original Primary will be converted to a secondary. Right-click on this secondary RVG and select **Start Replication** with **Synchronize Automatically** option.

- 3 After the sychronization is complete, perform a migrate operation to transfer the Primary role back to the original Primary. To do this, right-click on the Primary RVG and select **Migrate** option from the menu that appears.
- 4 Ensure that the volumes have retained the same drive letters that existed before the disaster.
- 5 Bring the RVG resource online on the Secondary.
- 6 Bring the EVS group online on the original Primary.

Index

C

- campus cluster
 - changing quorum resource to dynamic mirrored quorum resource
 - Windows Server 2008 353
 - configuration 272
 - creating SFW cluster disk groups and volumes 294
 - installing and configuring hardware 280, 331
 - making quorum cluster disk group a cluster resource 304, 352
 - Microsoft cluster configuration 269, 272, 323
 - Microsoft cluster failure scenarios 277, 329
 - prerequisites 269, 320
 - setting up a group for the application 303
 - verifying cluster configuration 314, 359
 - Vxclus 279, 331
- Cluster Administrator 191, 210, 386
- clustering concepts
 - ownership of quorum 278, 330
 - quorum 278, 330
- commands for snapshots
 - prepare 160
 - snap back 160
 - snap clear 155, 160, 161
 - vxsnap 169
 - vxsnap create 172
 - vxsnap prepare 171
 - vxsnap reattach 175
 - vxsnap restore 176
- configuration
 - Exchange storage for Quick Recovery 55
 - Microsoft cluster with SFW 188, 236
 - Microsoft cluster with SFW and VVR 368, 416

D

- disk groups
 - creating
 - Exchange Quick Recovery 56
 - Exchange with Microsoft cluster 56, 202, 249, 296
 - overview
 - Exchange with Microsoft cluster 201, 248
 - planning for Microsoft campus cluster 294
 - planning for Microsoft cluster 201, 248
 - quorum resource with mirrored volume 226, 257
- disk space requirements 187, 235, 271, 322, 367, 413
- drive letters 134, 141, 149
- driver signing options
 - resetting 381
- dynamic mirrored quorum
 - creating 225

E

- Exchange
 - installing for Microsoft cluster 208, 259
 - System Attendant 217
- Exchange database resource, dependency on the RVG resource 447
- Exchange databases
 - dismounting for hardware recovery 139, 147
 - moving to shared storage 221, 261
 - setting paths for Quick Recovery 60
- Exchange storage configuration for Quick Recovery 55
- Exchange transaction log paths for Quick Recovery 60
- Exchange virtual server group, creating for Microsoft cluster 210, 260, 386, 433

F

- failover clustering *see* Microsoft cluster
- FastResync 33
- First Storage Group 62, 67, 223
- FlashSnap 33
- forest and domain, preparing for Exchange 208

H

- hardware installation for cluster 280, 331
- hardware recovery
 - database and transaction logs volumes
 - missing 129
 - deleting missing volumes 133
 - identifying the missing volumes 131, 138
 - prerequisites 128
 - replacing hardware 133, 148
 - sample configuration 128
 - tasks 126
 - using snapshots 125

L

- local continuous replication (LCR) snapshots 50

M

- Microsoft cluster
 - creating an Exchange virtual server group 210, 260, 386, 433
 - creating dynamic quorum resource
 - Windows Server 2003 225
 - Windows Server 2008 256
 - establishing the cluster
 - Windows Server 2003 191, 282, 334
 - Windows Server 2008 239, 282
 - installing SFW 194, 242
 - process overview
 - Windows Server 2003 183
 - Windows Server 2008 231
 - Quick Recovery snapshot considerations 48
 - setting up a group for application 303
 - verifying configuration
 - Windows Server 2003 228
 - Windows Server 2008 263

Microsoft cluster with VVR

- bringing up Exchange on the secondary site 405, 450
- configuring the primary site 371, 419
- installing SFW 371, 419
- monitoring the status of the replication 404, 449
- normal operations and recovery
 - procedures 404, 449
- planned migration 404, 449
- process overview 361, 409
- replication recovery procedures 405, 450
- restoring the primary site 406, 451
- setting up the secondary site 382, 430
- system requirements 367, 415
- Microsoft Exchange Writer Replica 50
- mirrors
 - adding to volumes whose drive letters or mount points were reassigned 162
 - preparing from VEA or vxsnap 92
 - scheduling preparation with Quick Recovery Configuration Wizard 74
- missing volumes
 - deleting 133, 139, 148
 - identifying 146
- mount points
 - changing 134
 - snapshot volumes 141, 149
- MSCS *see* Microsoft cluster
- MSDTC resource 192

N

- network and storage, configuring
 - Microsoft cluster with SFW 190, 238

P

- passive copy snapshots 50, 113
- pre and post snapshot script file locations in VCS
 - cluster 48
- prepare command 160
- prerequisites
 - hardware recovery 128
 - Quick Recovery 54
 - see also* requirements

Q

- Quick Recovery
 - best practices 45
 - components 33
 - configuring Exchange storage 53
 - methods 41
 - overview 31
 - planning 37, 44
 - prerequisites 54
 - process 33
 - recommendations 45
 - recovery 111
 - storage requirements 39
 - system requirements 37
 - VCS considerations 47
 - VVR considerations 50
- Quick Recovery Configuration Wizard
 - overview 69
 - prerequisites 73
 - running 74
- quorum resource
 - changing to dynamic mirrored 225, 228, 258
 - creating cluster disk group 226, 257
 - with mirrored volume 226, 257

R

- recovery
 - adding disks to the dynamic disk group 140
 - after hardware failure 125, 179
 - passive copy snapshots 113
 - prerequisites 128
 - replacing hardware 140
 - single database 113
 - storage group 113
 - tasks for hardware recovery 126
 - VSS Restore Wizard 113
 - vxsnap utility 120
- recovery storage group (RSG) 112, 119
- Replicated Data Set (RDS), setting up 389, 435
- requirements
 - disk space 271, 322
 - Microsoft cluster with SFW 186, 233
 - Microsoft cluster with SFW and VVR 365, 412
 - Quick Recovery 37
- resetting
 - driver signing options 381
- resources
 - creating IP for VVR 387, 434

- roll-forward recovery
 - to point of failure 178
 - to point of failure of a single database 179
- RVG resource, creating (primary and secondary sites) 400, 446

S

- SA resource, dependency on the RVG resource
 - resources
 - dependencies 401
- schedules for snapshot sets
 - creating with Quick Recovery Configuration Wizard 74
 - creating with VSS Snapshot Scheduler Wizard 93
 - deleting 86, 87
 - modifying 86
 - synchronizing after adding a cluster node 87
 - troubleshooting 84
- SFW disk space requirements 187, 235, 367, 413
- SFW, installing with MSCS/Failover Cluster
 - option 194, 242
- snap back 160
- snap clear 155, 160, 161
- snapshot volumes
 - changing drive letter 134, 142, 150
 - changing mount points 134
 - reattaching healthy volumes 154, 160
- snapshots
 - clearing association for volumes 155, 161
 - creating after hardware recovery 156, 163
 - creating mirrors after hardware recovery 155, 162
 - LCR passive copy 50
 - methods of implementing 41
 - Microsoft cluster considerations 48
 - modifying schedules 86
 - overview 70
 - planning 37
 - reattaching split-mirror 108
 - recovery storage group (RSG) 112
 - refreshing after hardware recovery 137, 145, 152
 - refreshing manually 107
 - refreshing on the current disks 153
 - scheduling with VSS Snapshot Scheduler Wizard 93
 - storage requirements 39
 - system requirements 37

- templates 71
- troubleshooting 84
- VCS considerations 47
- viewing status 83
- VVR considerations 50
- Solutions Configuration Center
 - context sensitivity 23
 - overview 21
 - running wizards remotely 25
 - starting 22
 - wizard descriptions 25
- storage group
 - recovering to the point of failure 144, 151
 - restoring to the point in time 136, 143
- storage requirements for Quick Recovery 39
- synchronizing schedules in a cluster 87
- System Attendant resource
 - dependency on the RVG resource 401, 447

T

- templates for snapshot sets
 - description 71
 - multiple components 72
 - schedule start dates 72

V

- VCS cluster
 - Quick Recovery considerations 47
- Volume Manager Disk Group resource 259, 260, 356
- Volume Shadow Copy Service 34
- volumes
 - creating
 - Microsoft cluster 204, 251, 298, 346
 - Quick Recovery 57
 - creating on primary 298
 - overview
 - Microsoft cluster 201, 248
 - planning for campus cluster 295, 343
- VSS
 - defined 34
 - framework 34
- VSS Restore Wizard 113
- VSS Snapback Wizard 108
- VSS Snapshot Scheduler Wizard 93

VVR

- creating replicator log volumes 389, 435
- Quick Recovery considerations 50
- setting up the replicated data sets (RDS) 389, 435
- VVR resources, creating on the primary and secondary sites 387, 434
- Vxclus utility 279, 331
- vxsnap create 172
- vxsnap prepare 171
- vxsnap reattach 175
- vxsnap restore 176
- vxsnap utility command reference 169
- vxsnap utility, using for recovery 120

X

- XML file locations 47