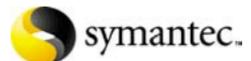


Veritas Storage Foundation™ and High Availability Solutions Solutions Guide

Windows Server 2003,
Windows Server 2008

5.1 Service Pack 2



Veritas Storage Foundation and High Availability Solutions Solutions Guide

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

Product version: 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 customer care_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

About the solutions guides	25
About Quick Recovery	26
About high availability	26
About campus clusters	26
About disaster recovery	26
About Microsoft clustering solutions	27
How this guide is organized	27

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	29
Starting the Configuration Center	30
Available options from the Configuration Center	31
About running the Configuration Center wizards	34
Following the workflow in the Configuration Center	36
Solutions wizard logs	37

Chapter 3 SFW best practices for storage

Best practices for storage availability	40
Best practices configuring SFW disk groups for availability	42
Best practices for storage performance	43
Best practices for I/O performance tuning	45
Best practices for storage capacity management	47

Section 2 Quick Recovery

Chapter 4 Quick Recovery overview

About the Quick Recovery solution	51
Need for implementing the SFW Quick Recovery solution	52
Understanding the underlying components of	

SFW's Quick Recovery process	53
FlashSnap	53
FastResync (FR)	54
Microsoft Volume Shadow Copy Service (VSS)	55
Overview of the Quick Recovery process	56
Creating initial snapshots	56
Refreshing a snapshot	56
Recovering a database	57
Other applications for point-in-time snapshots	58
Off-host backups	58
Reporting and analysis	59
Application testing and training	60

Chapter 5 Quick Recovery example

Example of Quick Recovery of an Oracle database	62
Create split-mirror snapshot of database	62
Recover database using split-mirror snapshot and database logs	63
More on FlashSnap: Tips and references	65

Chapter 6 Adding DMP to a clustering configuration

About dynamic multi-pathing	68
Overview of configuration tasks for adding DMP DSMs	69
Reviewing prerequisites	69
Reviewing the configuration	71
Steps for a new cluster configuration	72
Steps for an existing cluster configuration	72

Section 3 High Availability

Chapter 7 High availability: Overview

About high availability	75
About clusters	75

Chapter 8 Deploying SFW HA for high availability: New installation

About the high availability solution	78
Tasks for a new high availability (HA) installation—additional applications	78
Reviewing the requirements	80
Disk space requirements	80
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	80

Supported operating systems	81
System requirements for SFW HA	84
Network requirements for SFW HA	84
Permission requirements for SFW HA	85
Additional requirements for SFW HA	85
Best practices for SFW HA	86
Reviewing the configuration	86
Configuring the storage hardware and network	88
Installing Veritas Storage Foundation and	
High Availability for Windows	91
Setting Windows driver signing options	91
Installing Storage Foundation HA for Windows	93
Resetting the driver signing options	96
Configuring disk groups and volumes	97
Planning disk groups and volumes	97
Creating dynamic cluster disk groups	99
Creating dynamic volumes	101
Managing disk groups and volumes	104
Importing a disk group and mounting a volume	104
Unmounting a volume and deporting a disk group	105
Configuring the cluster	106
Configuring notification	117
Installing and configuring the application or server role	121
Configuring a File Share server role	121
Configuring a Print Share server role	121
Installing and configuring the IIS application	122
Installing and configuring Microsoft Virtual Server	123
Installing additional applications	123
Configuring the service group	124
Configuring the File Share service group	124
Configuring the PrintShare service group	134
Configuring the IIS service group	145
Configuring the MSVirtual Machine service group	154
Configuring the service group for any additional applications	157
Configuring Application Dependencies	174
Configuring an Oracle service group	177
Verifying the cluster configuration	181
Possible tasks after completing the configuration	182
Configuring the Cluster Management Console connection	183
Prerequisites for installing the cluster connector	183
Installing the cluster connector on Windows clusters	184
Configuring the cluster connector	185
Modifying the existing cluster configuration	188

Modifying the application service groups	193
Modifying the FileShare service group	194

Section 4 Campus Clustering

Chapter 9 Introduction to campus clustering

About Campus Clusters	206
Sample campus cluster configuration	206
Differences between campus clusters and local clusters	208

Chapter 10 Deploying SFW HA for campus cluster

About the Campus Cluster solution	210
Reviewing the requirements	212
Disk space requirements	212
Requirements for Veritas Storage Foundation	
High Availability for Windows (SFW HA)	212
Supported operating systems	213
System requirements for SFW HA	216
Network requirements for SFW HA	216
Permission requirements for SFW HA	217
Additional requirements for SFW HA	217
Best practices for SFW HA	218
Reviewing the configuration	219
Overview of campus clustering with VCS	221
Reinstating faulted hardware	222
Setting the ForceImport attribute	224
Installing and configuring the hardware	224
Configuring the storage hardware and network	225
Installing Veritas Storage Foundation HA for Windows	228
Setting Windows driver signing options	228
Installing Storage Foundation HA for Windows	230
Resetting the driver signing options	233
Configuring the cluster	233
Configuring notification	245
Creating disk groups and volumes	249
Configuring the disks and volumes	250
Creating a dynamic (cluster) disk group	251
Creating a volume	253
Installing the application on cluster nodes	257
Configuring a File Share server role	258
Configuring a Print Share server role	258
Installing and configuring the IIS application	258

Installing and configuring Microsoft Virtual Server	259
Installing additional applications	260
Exporting and importing a disk group	260
Configuring service groups	261
Configuring the File Share service group	263
Configuring the Print Share service group	273
Configuring the IIS service group	284
Configuring the MSVirtual Machine service group	293
Configuring the service group for any additional applications	296
Configuring Application Dependencies	313
Configuring an Oracle service group	316
Verifying the cluster configuration	320

Section 5 Replicated Data Clusters

Chapter 11 About Replicated Data Clusters

About Replicated Data Clusters	326
How VCS Replicated Data Clusters work	327
Setting up a Replicated Data Cluster configuration	328
Setting up replication	328
Configuring the service groups	329
Migrating the service group	331

Chapter 12 Deploying Replicated Data Clusters: New application installation

Tasks for a new replicated data cluster installation—additional applications	334
Reviewing the requirements	336
Disk space requirements	336
Requirements for Veritas Storage Foundation	
High Availability for Windows (SFW HA)	336
Supported operating systems	337
System requirements for SFW HA	340
Network requirements for SFW HA	340
Permission requirements for SFW HA	341
Additional requirements for SFW HA	341
Best practices for SFW HA	342
Reviewing the configuration	342
Sample configuration	343
Configuring the storage hardware and network	343
Installing Veritas Storage Foundation HA for Windows	346
Setting Windows driver signing options	346

Installing Storage Foundation HA for Windows	348
Resetting the driver signing options	352
Configuring VxSAS	352
Configuring the cluster	355
Configuring notification	366
Configuring disk groups and volumes	370
Planning disk groups and volumes	370
Creating dynamic cluster disk groups	372
Creating dynamic volumes	373
Managing the disk group and volumes	377
Importing a disk group and mounting a volume	377
Unmounting a volume and deporting a disk group	378
Installing and configuring the application or server role	379
Configuring a File Share server role	379
Configuring a Print Share server role	379
Installing and configuring the IIS application	380
Installing and configuring Microsoft Virtual Server	381
Installing additional applications	381
Configuring the service group	382
Configuring the File Share service group	382
Configuring the Print Share service group	392
Configuring the IIS service group	403
Configuring the MSVirtual Machine service group	411
Configuring the service group for any additional applications	414
Configuring Application Dependencies	430
Creating the primary system zone	433
Verifying the installation in the primary zone	434
Creating a parallel environment in the secondary zone	435
Adding the systems in the secondary zone to the cluster	436
Setting up the Replicated Data Sets (RDS)	441
Configuring a hybrid RVG service group for replication	452
Creating the RVG service group	452
Configuring the RVG service group for RDC replication	453
Configuring the RVG Primary resources	459
Configuring the primary system zone for the RVG	461
Setting a dependency between the service groups	462
Adding the nodes from the secondary zone to the RDC	463
Adding the nodes from the secondary zone to the RVG service group	463
Configuring secondary zone nodes in the RVG service group	464
Configuring the IP resources for failover	464
Adding the nodes from the secondary zone to the application service group	467

Configuring the zones in the application service group	468
Verifying the RDC configuration	469
Bringing the service group online	469
Switching online nodes	469
Additional instructions for GCO disaster recovery	470

Section 6 Disaster Recovery

Chapter 13 Disaster recovery: Overview

About a disaster recovery solution	474
Need for implementing a disaster recovery solution	475
Overview of the recovery process	476
Components of VVR that enable disaster recovery	477
Understanding replication	477
Modes of replication	477
Features of VVR that help in disaster recovery	478

Chapter 14 Deploying disaster recovery: New application installation

Tasks for a new disaster recovery installation—additional applications	482
Reviewing the requirements	485
Disk space requirements	485
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	486
Supported operating systems	486
System requirements for SFW HA	490
Network requirements for SFW HA	490
Permission requirements for SFW HA	491
Additional requirements for SFW HA	491
Best practices for SFW HA	492
Reviewing the configuration	492
Supported disaster recovery configurations for service group dependencies	493
Configuring the storage hardware and network	494
Managing disk groups and volumes	497
Importing a disk group and mounting a volume	497
Unmounting a volume and deporting a disk group	498
Setting up the secondary site: Configuring SFW HA and setting up a cluster	498
Installing SFW HA	499
Setting Windows driver signing options	499

Installing Storage Foundation HA for Windows	501
Resetting the driver signing options	504
Configuring the cluster	504
Configuring notification	516
Verifying that your application or server role is configured for HA at the primary site	520
Setting up your replication environment	520
Configuring the VVR security service	521
Requirements for EMC SRDF array-based hardware replication	523
Requirements for Hitachi TrueCopy array-based hardware replication	526
Assigning user privileges (secure clusters only)	528
Configuring disaster recovery with the DR wizard	529
Cloning the storage on the secondary site using the DR wizard (VVR replication option)	533
Creating temporary storage on the secondary site using the DR wizard (array-based replication)	537
Installing and configuring the application or server role (secondary site)	541
Installing the FileShare application	541
Installing the PrintShare application	541
Installing the IIS application	541
Installing the Microsoft Virtual Machine application	542
Installing additional applications	542
Cloning the service group configuration from the primary to the secondary site	543
Configuring replication and global clustering	546
Configuring VVR replication and global clustering	546
Configuring EMC SRDF replication and global clustering	554
Configuring Hitachi TrueCopy replication and global clustering	557
Configuring global clustering only	560
Verifying the disaster recovery configuration	563
Establishing secure communication within the global cluster (optional)	565
Adding multiple DR sites (optional)	567
Possible task after creating the DR environment:	
Adding a new failover node to a VVR environment	568
Preparing the new node	568
Preparing the existing DR environment	568
Modifying the replication and application service groups	569
Reversing replication direction	570
Maintaining: Normal operations and recovery procedures (VVR environment)	571

Normal operations: Monitoring the status of the replication	571
Performing planned migration	571
Disaster recovery procedures	572
Recovery procedures for service group dependencies	574

Chapter 15 Testing fault readiness by running a fire drill

About disaster recovery fire drills	577
About the Fire Drill Wizard	578
About Fire Drill Wizard general operations	578
About Fire Drill Wizard operations in a VVR environment	579
About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment	580
About post-fire drill scripts	582
Tasks for configuring and running fire drills	584
Prerequisites for a fire drill	586
Prerequisites for a fire drill in a VVR environment	586
Prerequisites for a fire drill in a Hitachi TrueCopy environment	587
Prerequisites for a fire drill in an EMC SRDF environment	588
Preparing the fire drill configuration	589
System Selection panel details	591
Service Group Selection panel details	591
Secondary System Selection panel details	591
Fire Drill Service Group Settings panel details	592
Disk Selection panel details	592
Hitachi TrueCopy Path Information panel details	593
HTCSnap Resource Configuration panel details	593
SRDFSnap Resource Configuration panel details	594
Fire Drill Preparation panel details	594
Running a fire drill	595
Recreating a fire drill configuration that has changed	597
Restoring the fire drill system to a prepared state	599
Deleting the fire drill configuration	600

Section 7 Microsoft Clustering Solutions

Chapter 16 Microsoft clustering solutions overview

About Microsoft clustering with high availability	605
About Microsoft clustering with Veritas Volume Replicator	606
About Microsoft clustering with campus clustering	606
About the SFW-Microsoft clustering-VVR configuration	606

Chapter 17 Deploying SFW with MSCS

Tasks for deploying SFW with MSCS (Windows Server 2003)	609
Reviewing the requirements	612
Supported software	612
Disk space requirements	614
System requirements	614
Reviewing the configuration	615
Configuring the network and storage	616
Establishing an MSCS cluster	618
Installing SFW	618
SFW installation tasks	619
Pre-installation tasks	619
Installing Veritas Storage Foundation for Windows	621
Post-installation tasks	624
Creating SFW disk groups and volumes	626
Planning disk groups and volumes	626
Creating dynamic cluster disk groups	628
Creating dynamic volumes	630
Setting up a group for the application in MSCS	633
Installing the application on cluster nodes	635
Completing the setup of the application group in MSCS	636
Implementing a dynamic quorum resource	637
Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume	638
Creating the quorum resource for the cluster group	638
Changing the quorum resource to a dynamic mirrored quorum resource	640
Verifying the cluster configuration	640

Chapter 18

Deploying SFW with Microsoft failover clustering

Tasks for deploying SFW with Microsoft failover clustering (Windows Server 2008)	644
Reviewing the requirements	646
Supported software for Microsoft failover clustering with SFW	646
Disk space requirements	647
System requirements	647
Reviewing the configuration	649
Configuring the storage hardware and network	650
Establishing a failover cluster	652
Installing SFW	653
SFW installation tasks	653
Pre-installation tasks	654
Installing Veritas Storage Foundation for Windows	654
Post-installation tasks	658

Creating SFW disk groups and volumes	659
Planning disk groups and volumes	659
Creating dynamic cluster disk groups	661
Creating dynamic volumes	663
Creating a group for the application in the failover cluster	667
Installing the application on cluster nodes	668
Completing the setup of the application group in the failover cluster	670
Implementing a dynamic quorum resource	671
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	671
Creating the quorum resource for the cluster group	673
Changing the quorum resource to a dynamic mirrored quorum resource	673
Verifying the cluster configuration	673

Chapter 19

Deploying SFW with MSCS in a campus cluster

Tasks for deploying SFW with MSCS in a campus cluster (Windows Server 2003)	675
Reviewing the requirements	678
Supported software	678
System requirements	679
Disk space requirements	680
Reviewing the configuration	681
Overview of campus clustering with MSCS	682
MSCS campus cluster failure scenarios	683
MSCS quorum and quorum arbitration	687
Configuring the network and storage	688
Establishing an MSCS cluster	690
Installing and configuring the operating system and MSCS on Server A	690
Configuring the shared storage and creating a partition for the Cluster quorum disk	691
Creating the first node of the cluster on Server A	691
Installing and configuring the operating system and MSCS on Server B	691
Connecting the two nodes	691
Creating the second node of the cluster on Server B	692
Verifying the cluster configuration	692
Installing SFW	693
SFW installation tasks	693
Pre-installation tasks	693
Installing Veritas Storage Foundation for Windows	695
Post-installation tasks	698

Creating disk groups and volumes	700
Configuring the disks and volumes	701
Creating a dynamic (cluster) disk group	702
Creating a volume	704
Changing the quorum resource to a dynamic quorum resource	708
Creating a dynamic cluster disk group for the quorum, mirrored	708
Making the quorum cluster disk group an MSCS resource	709
Changing the quorum resource to the dynamic mirrored quorum resource	711
Setting up a group for the application in MSCS	711
Installing the application on the cluster nodes	711
Completing the setup of the application group in MSCS	713
Verifying the cluster configuration	714

Chapter 20 Deploying SFW with Microsoft failover clustering in a campus cluster

Tasks for deploying SFW with Microsoft failover clustering in a campus cluster (Windows Server 2008)	718
Reviewing the requirements	720
Supported software	720
System requirements	721
Disk space requirements	722
Reviewing the configuration	723
Overview of campus clustering with Microsoft clustering	724
Campus cluster failure with Microsoft clustering scenarios	725
Microsoft clustering quorum and quorum arbitration	729
Configuring the network and storage	730
Establishing a Microsoft failover cluster	732
Connecting the two nodes	733
Installing SFW	735
SFW installation tasks	735
Pre-installation tasks	735
Installing Veritas Storage Foundation for Windows	736
Post-installation tasks	739
Creating disk groups and volumes	740
Configuring the disks and volumes	741
Creating a dynamic (cluster) disk group	743
Creating a volume	745
Implementing a dynamic quorum resource	750
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource	750
Adding the volume manager disk group for the quorum	750

Changing the quorum resource to the dynamic mirrored quorum resource	751
Setting up a group for the application in the failover cluster	752
Installing the application on the cluster nodes	754
Completing the setup of the application group in the cluster	755
Verifying the cluster configuration	757

Chapter 21 Deploying SFW and VVR with MSCS

Tasks for deploying SFW and VVR with MSCS (Windows Server 2003) ...	759
Part 1: Setting up the cluster on the primary site	763
Reviewing the prerequisites and the configuration	763
Installing and configuring the hardware	767
Installing Windows and configuring network settings	767
Establishing the cluster under MSCS (primary site)	769
Installing SFW (primary site)	769
Installing Veritas Volume Replicator Security Services (VxSAS)	770
Creating SFW disk groups and volumes	772
Setting up a group for the application in MSCS	773
Installing the application (Primary site)	775
Completing the setup of the application group in MSCS	776
Changing the quorum resource to a dynamic quorum resource	778
Testing of the cluster on the primary site	781
Part 2: Setting up the cluster on the secondary site	781
Repeating cluster configuration steps for the secondary site	783
Part 3: Adding the VVR components for replication	784
VVR components overview	784
Configuring the Replicator Log volumes for VVR	785
Setting up the Replicated Data Sets (RDS) for VVR	787
Creating an RVG resource and setting the dependencies	797
Part 4: Maintaining normal operations and recovery procedures	800
Normal operations: Monitoring the status of the replication	800
Performing planned migration	800
Disaster recovery procedures	801

Chapter 22 Deploying SFW and VVR with Microsoft failover clustering

Tasks for deploying SFW and VVR with Microsoft failover clustering (Windows Server 2008)	805
Part 1: Setting up the cluster on the primary site	810
Reviewing the prerequisites and the configuration	810
Supported software for Microsoft failover clusters with SFW	810
Installing and configuring the hardware	814
Installing Windows and configuring network settings	814

Establishing the Microsoft failover cluster (primary site)	814
Installing SFW (primary site)	816
Installing Veritas Volume Replicator Security Services (VxSAS)	816
Creating SFW disk groups and volumes	819
Completing the primary site configuration	820
Part 2: Setting up the cluster on the secondary site	821
Repeating cluster configuration steps for the secondary site	822
Part 3: Adding the VVR components for replication	823
VVR components overview	823
Configuring the Replicator Log volumes for VVR	824
Setting up the Replicated Data Sets (RDS) for VVR	826
Creating resources for VVR	836
Creating an RVG resource and setting the dependencies	837
Part 4: Maintaining normal operations and recovery procedures	840
Normal operations: Monitoring the status of the replication	840
Performing planned migration	840
Disaster recovery procedures	841

Section 8 Server Consolidation

Chapter 23 Server consolidation overview

Server consolidation definition	847
Need for implementing server consolidation	847
Advantages of using SFW with server consolidation	848
Overview of the server consolidation process	850

Chapter 24 Server consolidation configurations

Typical server consolidation configuration	852
Proof of concept	852
Server consolidation configuration 1—many to one	853
About this configuration	853
Reviewing the configuration requirements	856
Preparing to consolidate	857
Migrating the data to the large server	858
Adding the storage array	859
Completing the consolidation process	859
Server consolidation configuration 2—many	
to two: Adding clustering and DMP	860
About this configuration	860
Reviewing the configuration requirements	863
Adding the new hardware	864
Establishing the MSCS cluster	865

Adding SFW support to the cluster	865
Setting up MSCS cluster groups for the applications	866
Installing applications on the second computer	866
Completing the setup of the application group in MSCS	866
Changing the quorum resource to the dynamic quorum resource	867
Verifying the cluster configuration	867
Enabling DMP	867
SFW features that support server consolidation	868
Automatic volume growth	868
Features that support storage in a SAN	868
Performance monitoring	869
Server consolidation customer success story	869

Section 9 Appendix

Appendix A Deploying Disaster Recovery: Manual implementation

Reviewing the configuration	876
Part 1: Setting up the cluster on the primary site	879
Installing SFW HA	879
Configuring the VVR security service	879
Configuring the cluster	882
Configuring disk groups and volumes	899
Planning disk groups and volumes	900
Creating dynamic cluster disk groups	901
Creating dynamic volumes	903
Installing the application on cluster nodes	906
Creating VCS service groups	909
About VCS service groups	909
Service group example with a generic database application	909
Verifying the cluster configuration	919
Part 2: Setting up the parallel environment on the secondary site	921
Part 3: Adding the VVR components for replication	922
VVR components overview	922
Configuring the Replicator Log volumes for VVR	923
Setting up the replicated data sets (RDS) for VVR	925
Creating the VVR RVG Service group	936
Part 4: Adding GCO components for wide-area recovery	937
Prerequisites for a global cluster environment	937
Linking clusters by adding a remote cluster	938
Converting a local Service group to a global group	939
Additional global cluster administration tasks	942
Part 5: Maintaining: Normal Operations and recovery procedures	943

Normal operations: Monitoring the status of the replication	943
Performing planned migration	943
Disaster recovery procedures	944
Index	947

Introduction

This section contains the following chapters:

- [Introducing Veritas Storage Foundation and High Availability Solutions](#)
- [Using the Solutions Configuration Center](#)
- [SFW best practices for storage](#)

Introducing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [“About the solutions guides”](#) on page 25
- [“About Quick Recovery”](#) on page 26
- [“About high availability”](#) on page 26
- [“About campus clusters”](#) on page 26
- [“About disaster recovery”](#) on page 26
- [“About Microsoft clustering solutions”](#) on page 27
- [“How this guide is organized”](#) on page 27

About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions Solutions Guide* contains solutions for the following:

- Quick Recovery
- High availability (HA)
- Campus clusters
- Disaster recovery (DR)
- Microsoft clustering

Separate guides are available for Microsoft Exchange and Microsoft SQL solutions.

About Quick Recovery

Quick Recovery is the process of creating and maintaining on-host point-in-time images of dynamic volumes that can be used to quickly recover from logical errors in data files.

Quick Recovery is designed to augment your traditional backup methodology.

About high availability

The term high availability (HA) refers to a state where data and applications are highly available because software or hardware is in place to maintain the continued functioning in the event of computer failure. High availability can refer to any software or hardware that provides fault tolerance, but generally the term has become associated with clustering. Local clustering provides high availability through database and application failover. Veritas Storage Foundation HA for Windows (SFW HA) includes Veritas Storage Foundation and Veritas Cluster Server and provides the capability for local clustering.

About campus clusters

A campus cluster is a single cluster that stretches over two sites using fiber channel connectivity, with SAN connections for data mirroring and network connections for cluster communication. Although two sites are the most common, more than two can be used for additional redundancy.

Campus clusters provide disaster protection when an entire site goes down by locating the clustered servers in different buildings or areas. This solution provides a level of high availability that is above mirroring or clustering at a single site and is an alternative to using replication software.

About disaster recovery

Wide area disaster recovery (DR) provides the ultimate protection for data and applications in the event of a disaster. If a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away. Veritas Storage Foundation HA for Windows (SFW HA) provides the capability for implementing disaster recovery.

About 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 may be used with Veritas Storage Foundation for Windows to provide high availability for any application or server role.

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

How this guide is organized

Where applicable, the *Veritas Storage Foundation and High Availability Solutions Solutions Guide* is organized to follow the workflow in the Solutions Configuration Center.

See [Chapter 2, "Using the Solutions Configuration Center"](#).

When setting up a site for disaster recovery using the Configuration Center, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration. Likewise, in this guide, you first follow the instructions in the high availability section and then continue with the appropriate chapter in the disaster recovery section.

The Solutions Configuration Center includes a number of wizards that were not available in earlier versions of the product, including a Disaster Recovery wizard. The earlier methods of setting up disaster recovery manually, without the wizard, are available in an appendix section.

How this guide is organized

Using the Solutions Configuration Center

This chapter covers the following topics:

- [About the Solutions Configuration Center](#)
- [Starting the Configuration Center](#)
- [Available options from the Configuration Center](#)
- [About running the Configuration Center wizards](#)
- [Following the workflow in the Configuration Center](#)
- [Solutions wizard logs](#)

About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your Veritas Storage Foundation for Windows (SFW) or SFW High Availability (HA) environment. The Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2003, 2007, and 2010
- Microsoft SQL Server 2000, 2005, 2008, and 2008 R2
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server 2010 (high availability and disaster recovery solutions)
- Additional applications

Depending on the application, the following solutions may be available:

- High availability at a single site for a new installation

Starting the Configuration Center

- High availability at a single site for an existing server
- Campus cluster disaster recovery, including the following:
 - Campus cluster using Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003, 2007, 2010, and for Microsoft SQL Server 2005, 2008, and 2008 R2)
- Fire drill to test the fault readiness of a disaster recovery environment

The Solutions Configuration Center provides two ways to access Solutions wizards:

- The Applications tab lists solutions by application. It provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.
- The Solutions tab, for advanced users, lists wizards by solution without additional instructions.

Starting the Configuration Center

You can start the Configuration Center in the following ways:

- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

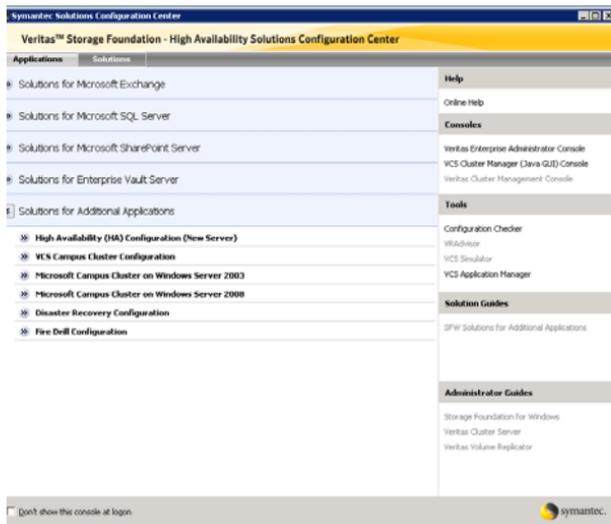
Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the solutions displayed when you click the application name are those available for that application. The Solution Guides listed in the right pane match the selected application.

In addition, some choices can vary depending on the operating system of the node on which you launch the wizard. For example, since Microsoft Exchange 2003 runs only on 32-bit operating systems, on a 64-bit system only the Exchange 2007 and Exchange 2010 configuration wizards are shown.

Figure 2-1 shows the choices available when you click Solutions for Additional Applications.

Figure 2-1 Solutions Configuration Center for additional applications



The submenu choices also vary by application. For example, different steps, information, or wizards are shown under High Availability (HA) Configuration for SQL Server than those shown for Additional Applications.

Figure 2-2 shows one of the steps for implementing high availability for SQL Server, showing the SQL Server configuration wizards.

Figure 2-2 Context-sensitive step for SQL Server



Figure 2-3 shows one of the steps for implementing high availability for additional applications, showing various configuration wizards.

Figure 2-3 Context-sensitive step for additional applications



About running the Configuration Center wizards

You can run the wizards from the Applications tab if you are walking through the configuration steps on the Solutions Configuration Center. If you are already familiar with configuration, you can also go directly to a particular wizard by selecting the Solutions tab.

The Configuration Center and some wizards can be run from a remote system. Wizards that you can run remotely include the following:

VCS Configuration Wizard	Sets up the VCS cluster
Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster Also can configure Veritas Volume Replicator (VVR) replication or configure the VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication Requires first configuring high availability on the primary site
Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots
Fire Drill Wizard	Sets up a fire drill to test disaster recovery Requires configuring disaster recovery first

Wizards related to storage configuration and application installation must be run locally on the system where the process is occurring. Wizards that you must run locally include the following:

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Microsoft Exchange Server 2003 and 2007 for the high availability environment If Microsoft Exchange is already installed, refer to the documentation for further instructions. This wizard is not required for Exchange Server 2010.
Exchange Configuration Wizard	Configures the service group for Microsoft Exchange Server 2003 and 2007 high availability
Exchange 2010 Configuration Wizard	Configures the service group for Microsoft Exchange Server 2010 high availability

SQL Server Configuration Wizard	Configures the service group for SQL Server 2000 or SQL Server 2005 high availability You must first install SQL Server on each node according to the instructions in the documentation.
SQL Server 2008 Configuration Wizard	Configures the service group for SQL Server 2008 or SQL Server 2008 R2 high availability You must first install SQL server on each node according to the instructions in the documentation.
Enterprise Vault Cluster Setup Wizard	Configures the service group for Enterprise Vault Server high availability
MSDTC Wizard	Configures an MSDTC server service group for SQL Server 2000, 2005, or 2008 environments
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group
SharePoint 2010 Configuration Wizard	Configures SharePoint Server 2010 service groups You can run the wizard from any SFW HA cluster node where SharePoint Server is installed and configured.

The Additional Applications section of the Configuration Center provides wizards to be run locally for creating service groups for the following applications or server roles:

File Share Configuration Wizard	Configures file shares for high availability
Print Share Configuration Wizard	Configures print shares for high availability
IIS Configuration Wizard	Configures IIS for high availability
MSVirtual Machine Configuration Wizard	Configures MS Virtual Machine for high availability
Oracle Agent Configuration Wizard	Configures Oracle for high availability
Application Configuration Wizard	Configures any other application service group for which application-specific wizards have not been provided

Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

Figure 2-4 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

Figure 2-4 Workflow for configuring high availability for additional applications



Solutions wizard logs

The Solutions Configuration Center provides access to many wizards. However, three wizards are built in to the Solutions Configuration Center:

- Disaster Recovery Wizard
- Fire Drill Wizard
- Quick Recovery Configuration Wizard

These three Solutions wizards are launched only from the Solutions Configuration Center, whereas other wizards can be launched from product consoles or the Start menu.

Logs created by these three Solutions wizards 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
```


SFW best practices for storage

This chapter includes the following topics

- [“Best practices for storage availability”](#) on page 40
- [“Best practices configuring SFW disk groups for availability”](#) on page 42
- [“Best practices for storage performance”](#) on page 43
- [“Best practices for I/O performance tuning”](#) on page 45
- [“Best practices for storage capacity management”](#) on page 47

Best practices for storage availability

For high availability of storage, use the following best practices to ensure continued access to data:

- [Adding software mirrors for critical data](#)
- [Locating data objects for optimum recovery](#)
- [Managing three-way software mirrors for reliability](#)
- [Software striping and mirroring on top of hardware RAID for high availability](#)

Adding software mirrors for critical data

For data that is absolutely critical to enterprise operation, use three-way mirrored volumes.

Using host-based volume management to construct the mirrors takes very little CPU (server) bandwidth, because the data and log writes are concurrent.

Additionally, connect the disks composing a three-way mirrored volume to the hosts using independent paths (cables, host bus adapters, connectors) to protect against path failure as well as disk failure.

Locating data objects for optimum recovery

When laying out volumes on disks with SFW, locate data objects that depend on each other on separate volumes, on separate disks. This ensures that a single disk failure does not destroy both data and its recovery mechanism.

Enterprise server environments often have interdependent sets of data. For example, the datasets in a database, its archive logs, and its redo log all depend on each other. If a volume holding database data fails, causing data loss, the typical practice would be the following steps:

- Repair the cause of the failure (for example, replace one or more disks).
- Restore the database to some baseline from a backup copy.
- Play the archive and redo logs against the restored copy to bring the database state as close to current as possible.

If the database logs reside on the same volume as the data, however, both data and logs are inaccessible, and database recovery is impossible.

Managing three-way software mirrors for reliability

The storage on a broken-off mirror should be restored to the original volume during periods of low application I/O load, if possible, because resynchronization and regeneration are I/O-intensive activities that can adversely affect application performance.

If a mirror is regularly broken off from a three-way mirrored volume, note that susceptibility to failure is greater during the interval between the third-mirror breakoff and the completion of resynchronization after the third-mirror storage is returned to the mirrored volume. By performing the restoration during periods of low I/O load, the susceptibility-to-failure window is minimized.

Software striping and mirroring on top of hardware RAID for high availability

For data protection in addition to the performance of striping, apply mirroring on top of striping. Perform the striping first and then mirror the striped volumes.

Striped volumes store large amounts of low-value or easily reproduced data where rapid access is required. Because striping alone will not maintain the availability of the data in a disk failure, consider mirroring also.

Best practices configuring SFW disk groups for availability

Veritas Storage Foundation for Windows supports multiple disk groups. Disk groups provide a way of organizing physical disks in a system into logical entities and simplify storage management for systems with large numbers of disks.

Disk groups are useful for managing storage in clusters, as well as convenient for organizing and managing disk storage resources for applications. SFW allows moving disks between host systems, providing an easy method of transferring storage from one system to another.

For high-availability with SFW disk groups, use the following best practices to ensure continued access to data:

- [Configuring disk groups for separate storage capacity pools or common pools](#)
- [Allocating disk groups for availability in clusters](#)

Configuring disk groups for separate storage capacity pools or common pools

Creating multiple disk groups creates separate storage capacity pools.

Effective use of subdisks is key to efficient disk group structure. The subdisks composing any given volume must be allocated from disks within a single disk group. Raw physical storage in one of these pools is available exclusively for use within the pool and cannot be used in other disk groups, unless an administrator specifically moves a disk from one disk group to another.

System administrators must decide on the basis of projected application and administrative needs whether to use disk groups to create disjointed storage pools or to manage all storage as a common pool.

Effective configuration of disk groups depends on an organization's application needs.

- If a critical application requires frequent volume expansion, allocating its storage in a private disk group helps guarantee that capacity is available when required. When storage capacity is added to the system, it is not absorbed by other applications.
- If a critical application unexpectedly requires additional storage and none is available in the disk group from which its volumes are allocated, the application will fail, even if the required amount of storage is available in other disk groups.

In general, multiple pools give the administrator greater flexibility, whereas a common pool may be more convenient for applications.

Allocating disk groups for availability in clusters

In a cluster, each application that fails over independently of other applications should have its data stored on volumes in disk groups exclusive to that application.

In a clustered environment, the disk group is the unit in which storage fails over from one computer to another. Only entire disk groups fail over. Thus, volumes that hold data for applications that are required to fail over should belong to disk groups that hold data for that application only. This allows an application's storage to fail over with the application and have no adverse effects on other applications or their associated storage. The disk groups should also be part of the application's resource group, so that failover can occur.

Best practices for storage performance

For optimal performance of storage, use the following best practices to ensure fast access to data:

- [Best practices for storage performance](#)
- [Host-based mirroring for increased read performance and failure tolerance](#)
- [Using software RAID 5 for read-mostly data](#)

Host-based mirroring for increased read performance and failure tolerance

Use host-based mirroring of virtual disks to increase overall system read performance and failure tolerance.

SFW provides host-based volume management to RAID subsystems, increasing overall data availability and I/O performance. With host-based volume management, software RAID can be applied across RAID subsystems from the same or different vendors, thus aggregating all the desirable properties of RAID subsystems.

In a mirrored configuration, read requests are handled in a round-robin fashion. The round-robin algorithm distributes read requests across all members, or "plexes," of a mirrored volume. Mirroring can increase read performance significantly.

Additionally, by configuring the hardware RAID subsystem-based virtual disks exported by different controllers as members of a host-based mirrored volume,

the host-based mirrored volume provides protection against I/O bus, host bus adapter, enclosure power and cooling, RAID controller, and disk failures.

Software striping across hardware for increased performance

Use SFW to combine multiple hardware arrays connected to the host via multiple buses in a single large striped volume, for higher transfer rates with some applications.

Host-based volume management can be used to aggregate the performance of multiple hardware subsystems by striping data across two or more virtual disks, each managed by a different RAID controller. Construct stripes across similar devices for the best use of storage. Because certain high-bandwidth applications, such as audio-visual streaming, have data transfer requirements that surpass the capability of a hardware array controller attached to the host by a single connection, the ability to aggregate the bandwidth of multiple data buses is needed.

Using software RAID 5 for read-mostly data

Host-based RAID-5 volumes are recommended for read-mostly data, because noticeable performance degradation may occur due to the overhead that writes generate.

Host-based RAID-5 volumes should be avoided in applications in which the rate of updates is high (more than about 10% of the aggregate I/O request-handling capacity of the disks constituting the volume), unless sufficient host CPU cycles are available. Disk controller RAID-5 volumes equipped with nonvolatile write-back cache may be used for more write-intensive applications (up to about 40% of the aggregate I/O request capacity of the disks composing the volume).

Best practices for I/O performance tuning

SFW enables administrators to “tune” any type of striped volume, including RAID-5 and mirrored striped volumes, by adjusting the stripe unit size. This feature is particularly useful for optimizing the I/O performance of these volume types.

Most I/O-bound applications can be characterized as one of the following:

- I/O-request intensive, making I/O requests faster than the hardware to which they are made can satisfy them
With rare exceptions, transaction-oriented applications (for example, credit verification, point of sale, order taking) are I/O-request intensive.
See [“Striping for I/O-request-intensive applications”](#) on page 45.
- Data-transfer intensive, moving large single streams of data between memory and storage
Scientific, engineering, audio, video, and imaging applications are typically data-transfer intensive.
See [“Striping for data-transfer-intensive applications”](#) on page 46.

If a striped volume will be used predominantly for one or the other of these I/O load types, the stripe unit size can be set at volume creation to optimize I/O performance.

Striping for I/O-request-intensive applications

A good compromise stripe unit size for I/O-request-intensive applications is one that results in a 3% to 5% probability of splitting in a uniform distribution of requests. For example, a 2 KB (four-block) database page size would have an ideal stripe unit size of 100 blocks. This would typically be rounded up to the nearest power of two (128 blocks, or 65,536 bytes) for simplicity.

I/O-request-intensive applications are typically characterized by small (for example, 2 to 16 KB) data transfers for each request. These applications are I/O bound because they make so many I/O requests, not because they transfer large amounts of data.

For example, an application that makes 1,000 I/O requests per second with an average request size of 2 KB uses at most 2 MB per second of data transfer bandwidth. Because each I/O request occupies a disk completely for the duration of its execution, the way to maximize I/O throughput for I/O-request-intensive applications is to maximize the number of disks that can be executing requests concurrently. Clearly, the largest number of concurrent I/O requests that can be executed on a volume is the number of disks that contribute to the volume’s storage. Each application I/O request that “splits” across two stripe units occupies two disks for the duration of its execution, reducing the number of

requests that can be executed concurrently and thus the efficiency of I/O response.

Therefore, try to minimize the probability that I/O requests “split” across stripe units in I/O-request-intensive applications.

The following factors influence whether an I/O request with a random starting address will split across two stripe units:

- The request starting address relative to the starting address of the storage allocation unit (the file extent)
- The size of the request relative to the stripe unit size

Most database management systems will allocate pages in alignment with the blocks in a file, so that requests for the first page will almost never split across stripe units. However, database requests for two or more consecutive pages may split across stripe units. In this case, larger stripe unit sizes reduce the probability of split I/O requests. However, the primary objective of striping data across a volume is to cause I/O requests to be spread across the volume’s disks. Too large a stripe unit size is likely to reduce this spreading effect.

Striping for data-transfer-intensive applications

The ideal stripe unit size for data-transfer-intensive applications that use a striped volume is the typical I/O request size of the application, divided by the number of data disks in the stripe. For example, if an application typically makes requests for 256 KB, an ideal stripe size for a four-disk striped volume would be 64 KB (256 KB/4).

Data-transfer-intensive applications typically request a large amount of data with every request, between 64 KB and 1 MB, or more. When a large amount of data is requested, the data-transfer phase of the request represents the majority of the request execution time. Thus, reducing data-transfer time improves I/O performance.

A single disk can transfer data only as fast as the data passes under the disk’s read-write head. For example, a disk that rotates at 10,000 RPM and has 200 blocks on a certain track cannot transfer data to or from that track any faster than 17.06 MB per second (200 blocks x 512 bytes per block/0.006 seconds per revolution). An application request for 500 KB would require five platter revolutions, or 30 milliseconds, to execute. If the request were addressed to a volume of five identical disks created with SFW, each disk would ideally deliver one-fifth of the data, and the request would complete in a shorter time.

In general, if a striped volume is optimized for data-transfer-intensive applications, each application I/O request will split evenly across all of the volume’s disks (or all but the disk containing parity data in the case of a RAID-5 volume).

Best practices for storage capacity management

Maintaining a percentage of unallocated storage capacity in a disk group is a useful means of managing online storage to avoid application failures. When an application requires more storage, its volumes can be extended quickly and easily by a system administrator while it's online, using the unallocated capacity. If volume expansion causes unallocated capacity to drop below a safety threshold, the event can be displayed in the GUI provided with SFW. Additional storage should then be installed and added to the disk group to maintain an adequate cushion for anticipated application requirements.

For storage capacity management, use the following best practices to ensure the best allocation of data:

- [Managing storage allocation for flexibility](#)
- [Aggregating hardware RAID for very large volumes](#)
- [Managing unallocated space for free space savings](#)
- [Reserving spares for failure-tolerant volume recovery](#)

Managing storage allocation for flexibility

One way to maximize the flexibility of storage allocation is to manage the disks in a disk group in units of a single capacity or of a small number of discrete capacities. This maximizes SFW's flexibility to allocate storage when new subdisks are required for new volumes, for volume extension, or for moving a subdisk from one disk to another.

To ensure that the amount of unallocated storage in each disk group is adequate, an appropriate level of unallocated storage be maintained. The distribution of unallocated storage across disks must allow for management operations such as failure-tolerant volume expansion to be carried out without violating volume failure tolerance and performance restrictions.

For example, if an additional mirror must be added to a mirrored striped volume, each subdisk of the added mirror must be located either on the same disk as the subdisk it extends, or on a disk separate from any of the volume's existing subdisks. (A subdisk is defined as a number of consecutively addressed blocks on a disk.) Subdisks are created by SFW as building blocks from which volumes are created. When an administrator makes a request to extend a volume, SFW checks the unallocated space in the disk group containing the volume to make sure that extension is possible. An administrator must maintain a distribution of unallocated capacity that allows such operations.

Aggregating hardware RAID for very large volumes

Combine LUNs from multiple RAID controllers with SFW to construct a very large volume capable of holding a very large database or file system, spanning multiple LUNs across controllers. This combined capability can give users better access to their data than if the file system or database is split across multiple LUNs.

For any type of hardware RAID used in an array, the size of a database or file system is limited to the maximum size of a logical unit number (LUN) in the particular hardware array. However, this limitation is removed with advanced host-based volume management.

Managing unallocated space for free space savings

Any policy for maintaining a minimum percentage of a disk group's capacity as unallocated space should include a cap to avoid maintaining wastefully large amounts of free space.

How much unallocated capacity to maintain depends strongly on application characteristics. In most cases, however, there are lower and upper bounds beyond which less or more unallocated storage would be of little use.

For example, an installation may observe a policy of maintaining a level of 8% to 10% of a disk group's total capacity as unallocated space. As the capacity of the disk group grows, however, the amount of unallocated space maintained by this policy can grow beyond any reasonable expectation of exploiting it effectively. If unallocated space is typically used in quantities of around 1 to 10 GB to relocate subdisks or to accommodate data processing peaks, then growing the disk group to 1 TB total capacity would mean that 100 GB are reserved for this purpose. If the typical number of subdisk moves or volume adds is one or two, a significant amount of storage capacity would never be used.

Reserving spares for failure-tolerant volume recovery

Reserve one or more spare disks for every 10 disks that are part of a failure-tolerant volume, with a minimum of one spare disk for any disk group that contains failure-tolerant volumes.

Storage capacity is managed in subdisk units, but entire disks usually fail. Because entire disks fail, spare capacity reserved for recovering from disk failures should be entire disks whose capacity is at least as large as that of the largest disk in a failure-tolerant volume in the disk group.

When a disk fails, all non-failure-tolerant volumes having subdisks on it fail, and all failure-tolerant volumes become degraded.

Quick Recovery

This section contains the following chapters:

- [Quick Recovery overview](#)
- [Quick Recovery example](#)

Quick Recovery overview

The chapter's topics are:

- [“About the Quick Recovery solution”](#) on page 51
- [“Need for implementing the SFW Quick Recovery solution”](#) on page 52
- [“Understanding the underlying components of SFW's Quick Recovery process”](#) on page 53
- [“Overview of the Quick Recovery process”](#) on page 56
- [“Other applications for point-in-time snapshots”](#) on page 58

About the Quick Recovery solution

Veritas FlashSnap is an option to SFW that is a highly efficient procedure involving multiple commands that allows detaching of a mirrored volume. Once the volume is detached, it can be used for a variety of purposes. This chapter focuses on the SFW Quick Recovery solution, which uses split-mirror snapshots to recover from logical errors in data files. It also gives a summary of the other uses for split-mirror snapshots.

Quick Recovery is the process of using on-host point-in-time copies of production data and a transaction log to recover a database that has been corrupted or that has missing data. If a database becomes corrupted, for example, you could reload the original data from the most recent snapshot, and then use the transaction log to bring the database current to the point before the corruption.

The SFW Quick Recovery solution uses on-host, disk-based snapshots to provide fast recovery from logical errors and eliminates the time-consuming process of restoring data from tape.

If you are using Microsoft Exchange, SFW has recovery procedures for Microsoft Exchange storage groups or individual databases within an Exchange storage group. Additionally, Quick Recovery of Microsoft SQL 2005 databases is

supported. Those procedures are provided through SFW's `vxsnap restore` command and the VSS Snapshot wizards.

SFW also provides the `vxsnapsql` utility for Quick Recovery which integrates with VDI to perform snapshot operations on SQL Server database volumes while the database is online and available. VDI quiesces the database for the short period of time required to create the snapshot and then immediately thaws it. This quiescing allows SQL snapshots to be taken while the database application remains active. The `vxsnapsql` utility is supported for SQL 2000 and SQL 2005. This chapter gives a general overview of SFW's Quick Recovery solution. For detailed information about implementing a Quick Recovery solution with SFW and Microsoft Exchange Server, see the *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange*.

For detailed information about implementing a Quick Recovery solution with SFW and Microsoft SQL Server, see the *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*.

For information about implementing Quick Recovery solution with Oracle database, see "[Example of Quick Recovery of an Oracle database](#)" on page 62.

Need for implementing the SFW Quick Recovery solution

Advantages of using SFW's Quick Recovery solution:

- **Faster than Restoring from Tape or Other Media**
On-host snapshot recovery is faster than restoring a full backup from tape or other media; this reduces downtime and helps meet service-level agreements for application availability. A Quick Recovery solution serves as a first line of defense to recover a corrupted database or missing data. The impact on system performance of maintaining a Quick Recovery image is limited to the brief time of detaching a split-mirror snapshot from its original volume.
- **A Less Costly and More Flexible Solution than Array-based Snapshots**
SFW's split-mirror snapshots are based on the FlashSnap technology. FlashSnap puts the snapshot logic on the host system itself, so you can use any storage you have or might acquire to create snapshots. This is in contrast to a split-mirror image created through a hardware storage array, where you are limited to only the storage provided by the array vendor. FlashSnap provides several benefits over a hardware-based approach:

- You can use virtually any storage hardware to create snapshots, including expensive arrays and simple JBOD storage.
- The volumes that are copied can span physical devices.
- The original and snapshot volumes can use different vendors' storage arrays.
- Integration with Microsoft Server Volume Shadow Copy Service
SFW integrates with the Windows Server Volume Shadow Copy Service (VSS) as both a VSS requestor and a VSS provider. This integration is provided by FlashSnap and SFW's `vxsnap` command line utility and VSS Snapshot wizards. The VSS process enables a VSS-aware application, such as Exchange, to be quiesced before the snapshot operation occurs and then resumed immediately after it. This pause of the application can produce Microsoft supported and guaranteed snapshots of your data. It protects the integrity of your data.
- Allows Multiple Snapshots at One Time
SFW offers the option to create simultaneous, multiple split-mirror snapshots. These snapshots can be done either through the GUI **Snap Shot** command or through the `vxsnap` CLI command.

Understanding the underlying components of SFW's Quick Recovery process

SFW's Quick Recovery solution uses Veritas FlashSnap and FastResync technology to leverage the Microsoft Volume Shadow Copy Service (VSS) capability to pause and resume a VSS-aware application.

FlashSnap

FlashSnap provides the ability to create and maintain the on-host point-in-time copies that are integral to the Quick Recovery solution. FlashSnap is the multi-step process used to create and maintain split-mirror snapshots that are copies of the original volumes they mirror. Both the original and snapshot volumes 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.

FlashSnap includes the following commands:

Prepare

Creates a snapshot mirror and attaches it to the original volume. The **Prepare** procedure may take considerable time because it involves creating a mirror, but it has to be done only the first time you perform the snap commands sequence.

Note: The **Prepare** command replaces the **Snap Start** command in the GUI. Both `prepare` and `snapstart` keywords are available in the CLI, however `prepare` is the recommended keyword.

Snap Shot

Detaches the snapshot mirror from the original volume. This split-mirror snapshot volume is an exact duplicate of the original database volume at the point in time the `snapshot` command is executed.

Snap Back

Reattaches the snapshot mirror to the original volume. The volumes can be resynchronized using either the original volume or the snapshot volume as the source. If a logical error has occurred on the original database volume, the snapshot volume can be used to quickly restore a consistent, point-in-time image to the original volume.

Snap Clear

Permanently removes the association between the snapshot volume and the original volume.

Snap Abort

Aborts the snapshot operation after a **Prepare** or **Snap Back** command is issued. **Snap Abort** permanently removes the snapshot mirror from the volume and releases its space.

The FlashSnap commands listed above are implemented through the SFW GUI. There are also command line equivalents, using the `vxassist` or `vxnsnap` command.

FastResync (FR)

The FastResync capability 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. When the snapshot volume is resynchronized with the original volume by using the **Snap Back** command, 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 Quick Recovery image can be refreshed with minimal impact on production.

FR is automatically enabled for a volume when the `prepare` operation is performed on the volume through the GUI **Prepare** command or the command line interface `vxassist snapstart` command.

Microsoft Volume Shadow Copy Service (VSS)

Microsoft Volume Shadow Copy Service (VSS) is a Windows service that provides the capability of creating snapshots or volume shadow copies. A volume shadow copy is a volume that represents a duplicate of the state of the original volume at the time the copy began. SFW integrates VSS into its snapshot function through the `vxsnap` command. Because SFW is a VSS requestor, it can initiate VSS snapshots at any time.

In the Windows Server 2003 version, `vxsnap` makes use of both FlashSnap and VSS technology to create high-quality snapshots that can be done when application files are open. VSS can quiesce the application for the moment when the snapshot is created and then resume the application immediately after the snapshot; but a VSS-aware application must be used, such as Microsoft Exchange Server.

For more information on how VSS and SFW work together, see the *Veritas Storage Foundation Administrator's Guide*.

Overview of the Quick Recovery process

The Quick Recovery process can be broken down into three phases: creating, refreshing, and recovering.

Creating initial snapshots

Split-mirror snapshots should be created on a regular schedule, following the backup of the database from tape. You can snapshot a database volume by itself or you can use the SFW GUI **Snap Shot** command or the `vxsnap` utility to snapshot one or more database volumes and any database log volumes simultaneously. If you have an application that is VSS-aware, such as Microsoft Exchange Server or Microsoft SQL Server 2005, you have the advantage of creating VSS snapshots. By taking VSS-enabled snapshots, you can create snapshot images without needing to take the database offline. Additionally, SFW offers the `vxsnapsql` utility for Microsoft SQL 2000 or Microsoft SQL 2005.

Creating a snapshot is a two-step process. The first step, **Prepare**, creates the snapshot mirror attached to the original volume. The second step, **Snap Shot**, detaches the snapshot mirror from the original volume and creates a separate on-host split-mirror snapshot volume.

Once a snapshot has been created, it can be refreshed quickly without repeating the time-consuming **Prepare** step.

Refreshing a snapshot

Periodically refresh or update your snapshot or set of snapshots so they contain a current copy of the original volumes. Refreshing a snapshot is a two-step process. During the first step, the **Snap Back** operation reattaches a snapshot volume to its original volume and uses Fast Resync to automatically update the snapshot mirror and synchronize it with the original volume, applying only the changes tracked in the Disk Change Object (DCO) volume. This process takes less time than the traditional method of copying the entire original volume to the returning mirror. In the second step, the **Snap Shot** operation is performed to detach the snapshot mirrors again, creating a new point-in-time copy of the database. If you are creating multiple snapshots, the SFW GUI **Snap Shot** command or the `vxsnap` CLI command must be used to snapshot all the database and log volumes simultaneously. This step is done without taking the database offline.

The **Snap Back** and `vxsnap` commands can be called from either the **bpnd_notify.bat** file in Veritas NetBackup or from a batch file in a pre/post command to run at the completion of a Veritas Backup Exec for Windows Servers backup job. Additionally, a script could be written and used with the

Windows Task Scheduler to automatically update the snapshot or set of snapshots on a regular basis.

Recovering a database

In the event a database needs to be recovered, you can use the snapshot or set of snapshots to restore the data.

Caution: Data corruption can occur if the FlashSnap utility does not have exclusive access to the volumes accessed in the **Snap Back** command. Before running the **Snap Back** command when using the snapshot data as the source, close any Explorer windows, applications, consoles, or third-party system management tools that may be accessing the volumes.

Veritas Storage Foundation for Windows provides recovery support for Microsoft Exchange storage groups or individual databases within an Exchange storage group. Through SFW's `vxsnap restore` command or the VSS Restore wizard, the VSS hot snapshots can be used for a point-in-time recovery of the storage group or a roll-forward recovery to the point of failure of either the storage group or an individual database within it.

Refer to the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange* for detailed procedures on using FlashSnap with Microsoft Exchange to perform hot snapshots and to implement recovery procedures.

For Microsoft SQL, you can use the snapshot volumes in a snapshot set to restore a corrupt database. You can restore a database to a specified point in time, the point of failure, or the point in time that the snapshot set was created (or last refreshed).

Refer to the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft SQL* for detailed procedures on Quick Recovery in a Microsoft SQL environment.

Other applications for point-in-time snapshots

This section describes several of the possible applications for using FlashSnap's snapshots for off-host processing. Topics include:

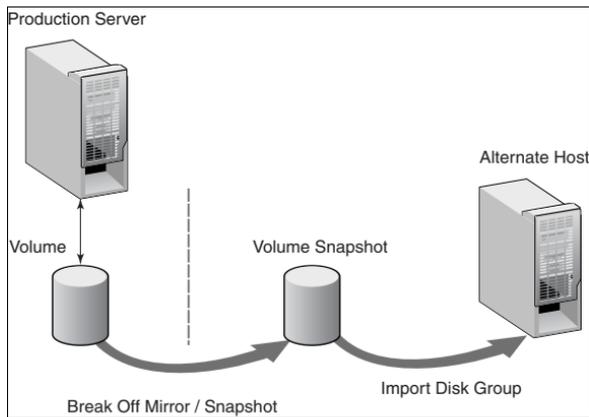
- “Off-host backups” on page 58
- “Reporting and analysis” on page 59
- “Application testing and training” on page 60

Off-host backups

The more frequent your backups, the less data lost or, in the case of a database with a transaction log, the faster your recovery. Incremental backups reduce the backup time but increase recovery time. For organizations with little or no backup window, off-host backups offer a good solution, particularly as the amount of data to be managed grows.

Because backups take place on another host, the backup window is of less concern, and you can make full backups each day. This speeds recovery time in the event a problem does occur.

Figure 4-1 Mirror break-off and import of the snapshot to the alternate host



FlashSnap simplifies the process of making snapshot volumes available for off-host processing with the Disk Group Split and Join feature. Using this feature,

administrators can split one or more volume snapshots into another disk group, then “deport” the disk group. The alternate host, running Storage Foundation for Windows, can then import that disk group and its volumes for off-host processing.

When the off-host processing is complete, you can rejoin the snapshot volume and its disk group in a similar manner, deporting it from the secondary host, importing it to the primary host, and rejoining the original disk group.

FlashSnap snapshots can be backed up with Veritas NetBackup or Backup Exec.

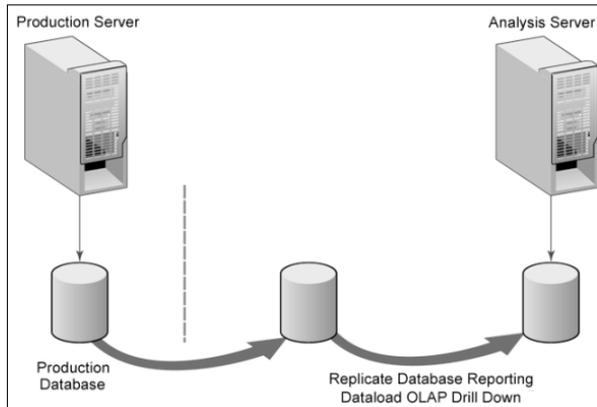
Reporting and analysis

Decision support and business intelligence are data-intensive activities that are critical to many organizations. Analysts and others frequently need to access up-to-date or even real-time data in their analysis. Retail organizations, for example, want to spot sales trends as they occur, and typically require at least daily updates. Financial institutions likewise must keep a close eye on current transactions to spot trends or potential problems quickly.

Unfortunately, reporting and analysis data needs typically conflict with the performance requirements of transactional database applications. Reporting and analysis activities generally implement a few selected statements that scan a large number of records, and may include complex processing. This will have an impact on the many simpler write and update activities characteristic of a transactional system. For this reason, among others, many companies load data from operational systems into data warehouses specifically designed and tuned for analytic queries. But even the process of creating the data loads can have a performance impact on your operational systems, causing most organizations to schedule these Extract, Transform, and Load (ETL) processes during off-hours, such as in the middle of the night.

You can solve this problem by creating point-in-time snapshots of the production systems to be used for reporting and analysis purposes. You can either run reports directly against the snapshot volumes or use the snapshots to extract data for a data load to the warehouse.

Figure 4-2 Extract, transform, and load (ETL) process



Because taking the snapshot itself has a very brief, limited impact on the production system, you can generate fresh data for analysis on a regular basis. You can even create a replica of the production database on a secondary system to be accessible for “drill-down” analysis from OLAP applications. Again, in the off-host scenario, the analysis has no impact on the production system.

Application testing and training

Software testing and training are other valuable applications for FlashSnap point-in-time copies. These are needs that cannot be addressed by simple data replication, because you need to be able to update and modify the copy of the data used for testing. FlashSnap addresses these needs easily.

By taking a snapshot and loading it on a host used for testing or development, you can provide developers and QA staff with the most realistic test data possible. By actually using a point-in-time copy of the production data, you can anticipate the behavior of the application in the production setting. You also save the time of creating and maintaining test data sets. This data can also be used for training purposes.

Quick Recovery example

This chapter provides a Quick Recovery example using an Oracle database.

The chapter's topics are:

- [“Example of Quick Recovery of an Oracle database”](#) on page 62
- [“More on FlashSnap: Tips and references”](#) on page 65

Example of Quick Recovery of an Oracle database

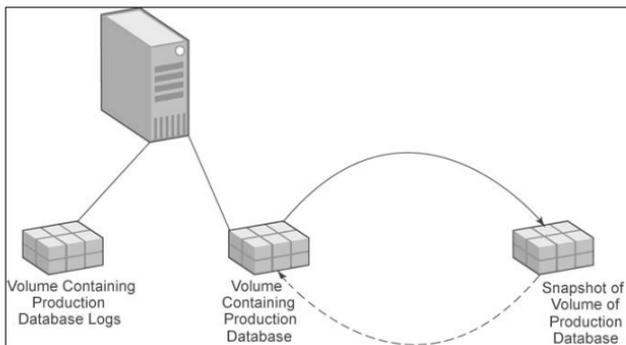
This example demonstrates how SFW's split-mirror snapshot can be used to recover an Oracle database after its data has become corrupted. The advantage of using the snapshot process is that it is much faster than recovering the database from tape backup. The process assumes that these split-mirror snapshots would take place on a regular schedule following the regular backup of the database.

This example does not require Microsoft Volume Shadow Copy Service (VSS).

Create split-mirror snapshot of database

The illustration below shows the snapshot step. The arrow pointing back to the original volume indicates that the snapshot volume can be rejoined to the original volume, updated, and ready to create a refreshed snapshot.

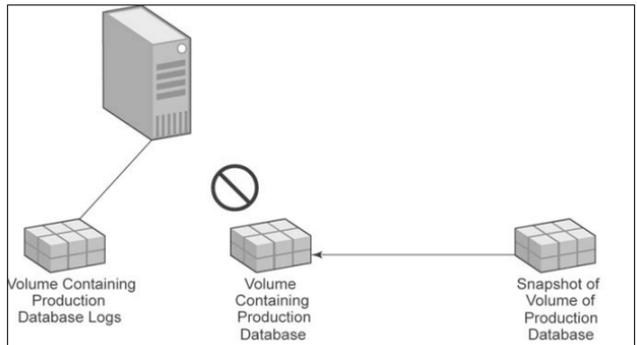
Figure 5-1 Creating a backup of the database with a snapshot



Recover database using split-mirror snapshot and database logs

The following illustration shows the situation where there has been a database failure. The snapshot volume is located on the right. The arrow pointing back to the production database volume represents the recovery of the database using the snapshot and applying the logs to bring the database to the level just before the failure occurred.

Figure 5-2 Database recovery from a snapshot



Overview of tasks

The main tasks for this example are:

- Snapshot the Oracle datafile volume.
- Resume normal processing with the Oracle datafile.
- Simulate Oracle datafile corruption.
- Recover the Oracle datafile.

Specific steps

Prerequisites

This example assumes:

- Experience in Oracle database backup and recovery.
- Experience with the SFW FlashSnap procedures.

Setup

- The Oracle database must be running on a single server.
- The Oracle database must be in ARCHIVELOG mode.
- The volume that contains the datafile for the Oracle database must meet the following requirements:
 - The volume must not be the system/boot volume.
 - The volume must be a SFW volume.

To snapshot the Oracle datafile volume

- 1 Open the Oracle database and verify that the tablespace you want to work with is running normally.
- 2 In SFW, prepare the volume that contains the datafile of the tablespace. The CLI command is:

```
vxassist -g<DynamicDiskGroupName> snapstart  
<DriveLetter>
```
- 3 In Oracle, ALTER the tablespace with the BEGIN BACKUP option to prepare the database logs for backup creation mode.
- 4 In SFW, use **Snap Shot** to make a snapshot of the Oracle datafile volume. The CLI command is:

```
vxassist -g<DynamicDiskGroupName> snapshot  
<DriveLetter>
```
- 5 In Oracle, ALTER the tablespace with the END BACKUP option to set the database logs to normal mode.
- 6 In Oracle, archive the current database log to keep it at the same level as the snapshot.

This completes the process of implementing the snapshot of the database and saving it for later use. This process offers the most benefit if it is done on a periodic basis.

Resuming normal processing with the Oracle datafile

- ◆ In Oracle, update the tables in the tablespace to create database log activity.

To simulate Oracle datafile corruption

- 1 In Oracle Enterprise Manager, offline the tablespace.
- 2 Use Windows Explorer to locate and open the datafile volume.
- 3 Delete the datafile.

To recover the Oracle datafile

- 1 In Oracle, take the datafile offline.
- 2 Use **Snap Back** in the VEA GUI to reattach the snapshot volume. Use the **Resynchronize using the snapshot** option.
- 3 In Oracle, use RECOVER TABLESPACE to apply the database logs to bring the replica to the level just before the datafile corruption occurred.
- 4 In Oracle, bring the datafile online.
- 5 Verify that the tablespace in the datafile has been recovered.

Note: This example uses a single snapshot of the datafile of the tablespace of an Oracle database. It could also be done by using multiple, simultaneous snapshots that include both the data file and the log.

More on FlashSnap: Tips and references

The following FlashSnap tips may be helpful:

- Use related disk group names.
For example, when doing off-host processing from the GUI, use “database” for the original disk group name and “database_snap” for the snapshot disk group name.
A disk group name can be a maximum of 18 characters long.
- When using FlashSnap with a database application, store all database files and related transaction logs on disks contained within a single dynamic disk group.
- For easy identification, the volumes within a disk group should begin with the name of the disk group. For example:
 DiskGroup1_VolumeName1
 DiskGroup1_VolumeName2
 DiskGroup1_VolumeName3
- For more information on FlashSnap, see the *Veritas Storage Foundation Administrator’s Guide*.
- For more information on Quick Recovery, see the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft Exchange* and the *Veritas Storage Foundation and High Availability Solutions, Quick Recovery and MSCS Solutions Guide for Microsoft SQL*.

Adding DMP to a clustering configuration

This chapter has the following topics:

- [“About dynamic multi-pathing”](#) on page 68
- [“Overview of configuration tasks for adding DMP DSMs”](#) on page 69
- [“Reviewing prerequisites”](#) on page 69
- [“Reviewing the configuration”](#) on page 71
- [“Steps for a new cluster configuration”](#) on page 72
- [“Steps for an existing cluster configuration”](#) on page 72

About dynamic multi-pathing

Dynamic Multi-pathing is an optional software component in SFW that provides redundant path support for your storage. This support is provided by DMP DSMs (DMP Device Specific Modules). You can add Dynamic Multi-pathing (DMP) to a clustering configuration.

The steps for adding Dynamic Multi-pathing are given for both a new cluster configuration and an existing cluster configuration. The cluster configuration can be either VCS or Microsoft clustering. The steps for enabling Dynamic Multi-pathing are done after a cluster is up and running and tested. A second host bus adapter in each computer allows redundant paths to the storage for fault tolerance purposes. The Dynamic Multi-pathing software controls the usage of the paths and allows only one path to the storage to operate at a time. However, if one path fails, the Dynamic Multi-pathing software will automatically transfer the storage to the second path.

In installing DMP DSMs you must connect only one host bus adapter path while you are setting up the hardware and installing the DMP DSM software and any other software. This is true for both a new installation and an existing installation.

For the hardware setup step for each server, install the second host bus adapter in each computer, but do not connect it to the switch. After SFW and the cluster are set up and working and the DMP DSMs software is installed and running, then perform the additional steps to activate the DMP DSMs.

Overview of configuration tasks for adding DMP DSMs

There are multiple tasks to add DMP DSMs on each server. The order in which the tasks are done in relation to the rest of the configuration depends on whether you are installing a new configuration or upgrading an existing cluster configuration. Detailed steps are presented in later sections. In summary, the tasks are the following:

- Install a second host bus adapter in each server. Do not connect the second path to each additional switch at this point. The path must be left unconnected.
For a new install, this step can be done with the initial hardware configuration before the servers are running.
For an existing cluster system, a rolling upgrade procedure is used to allow installation of the hardware and software on the inactive node on the cluster. When the installation of one node is complete, switch the active node and complete the hardware and software installation on the remaining node that is inactive.
- Install the DMP DSMs software on the inactive node.
Installing DMP DSMs requires a reboot, and this avoids rebooting the active node of the cluster.
- Using appropriate cables, connect the second path on Server A to the second switch. Configure the switch, if necessary. Do the same for the second switch on Server B.
- Verify that both paths are under DMP DSMs control.
- Access the Array Settings dialog for each array and make sure that the array load balancing settings are set to active/passive.

Reviewing prerequisites

This solution assumes that the required software is already installed and configured. Refer to the *Veritas Storage Foundation Administrator's Guide* and *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation and configuration information.

Supported software

Veritas Storage Foundation 5.1 for Windows (SFW) or Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Dynamic Multi-pathing (DMP) option.

If you already have the SFW or SFW HA software installed and want to add Dynamic Multi-pathing, purchase a license key for the Dynamic Multi-pathing option and use **Add or Remove Programs** from the Windows Control Panel to add the option.

Hardware requirements (Two-server cluster)

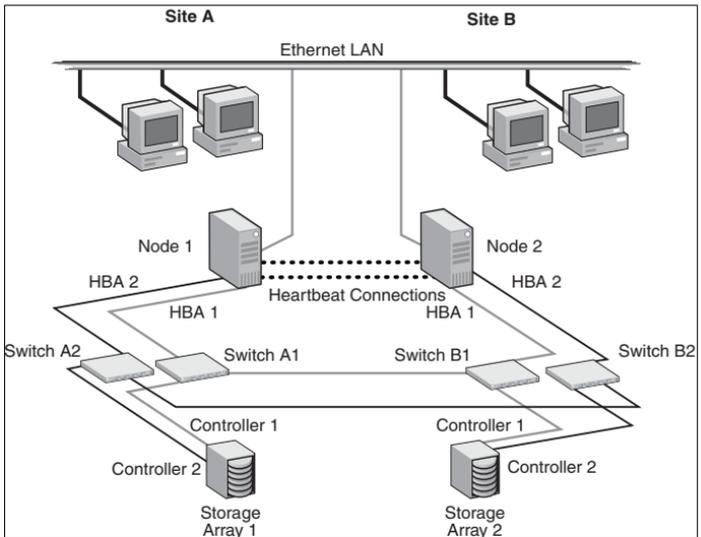
- 2 HBAs with appropriate cabling
- 2 fibre switches

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at: <http://www.symantec.com/business/support/index.jsp>.

Reviewing the configuration

The purpose of the configuration is to add Dynamic Multi-pathing (DMP) to a clustering configuration. Dynamic Multi-pathing is an optional software component in SFW and SFW HA that provides redundant path support for your storage. You need to have two switches per server, one switch to accommodate each path. The diagram illustrates the four different host bus adapter paths.

Figure 6-1 Adding fault tolerance with DMP



Steps for a new cluster configuration

Use the following procedure to set up DMP DSMs.

To set up DMP DSMs

- 1 Install additional hardware and its appropriate drivers.
- 2 Connect only one path from the array to the computer.
- 3 Install the appropriate DMP DSMs using **Add or Remove Programs** from the Windows Control Panel. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.
- 4 Reconnect the additional physical path.
- 5 Reboot the system.
- 6 Verify that the additional path is displayed.

Steps for an existing cluster configuration

Use the following procedure to set up DMP DSMs.

To set up DMP DSMs

- 1 Move resources to another node or take the resources offline.
- 2 Install additional hardware and its appropriate drivers.
- 3 Connect only one path from the array to the computer.
- 4 Install the appropriate DMP DSM using **Add or Remove Programs** from the Windows Control Panel. For more information, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.
- 5 Reconnect the additional physical path.
- 6 Reboot the system.
- 7 Verify that the additional path is displayed.

High Availability

This section contains the following chapters:

- [High availability: Overview](#)
- [Deploying SFW HA for high availability: New installation](#)
- [Adding DMP to a clustering configuration](#)

High availability: Overview

This chapter contains the following topics:

- [“About high availability”](#) on page 75
- [“About clusters”](#) on page 75

About high availability

A high availability solution maintains continued functioning of applications in the event of computer failure, where data and applications are available using redundant software and hardware. “High availability” can refer to any software or hardware that provides fault tolerance, but generally it has become associated with clustering. This section will focus on local clustering configurations that use Veritas Cluster Server (VCS) with Veritas Storage Foundation for Windows.

About clusters

A cluster is a group of independent computers working together as a single system to ensure that mission-critical applications and resources are highly available. The cluster is managed as a single system, shares a common namespace, and is specifically designed to tolerate component failures and to support the addition or removal of components in a way that is transparent to users.

Keeping data and applications functioning 24 hours day and seven days a week is a requirement for critical applications today. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

Deploying SFW HA for high availability: New installation

This chapter covers the following topics:

- [“About the high availability solution”](#) on page 78
- [“Tasks for a new high availability \(HA\) installation—additional applications”](#) on page 78
- [“Reviewing the requirements”](#) on page 80
- [“Reviewing the configuration”](#) on page 86
- [“Configuring the storage hardware and network”](#) on page 88
- [“Installing Veritas Storage Foundation and High Availability for Windows”](#) on page 91
- [“Configuring disk groups and volumes”](#) on page 97
- [“Configuring the cluster”](#) on page 106
- [“Installing and configuring the application or server role”](#) on page 121
- [“Configuring the service group”](#) on page 124
- [“Verifying the cluster configuration”](#) on page 181
- [“Possible tasks after completing the configuration”](#) on page 182

About the high availability solution

This chapter provides the steps for setting up a High Availability (HA) solution, using SFW HA in a new installation. The chapter describes the process for any generic application or server role and specifically for File Share, PrintShare, IIS and MSVirtual Machines.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA solutions for the example applications or server roles.

See “[Using the Solutions Configuration Center](#)” on page 29.

For examples of the SFW HA solution with Microsoft Exchange or Microsoft SQL Server, see the other Solutions Guides included with this release:

- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange*
- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*

Tasks for a new high availability (HA) installation—additional applications

This chapter provides information on how to install and configure the high availability and application components.

- | | |
|----------------|---|
| Active-Passive | One application instance per node with one to one failover capabilities. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. |
| Active-Active | Multiple application instances per cluster node.
For example, in a two-node cluster with two application instances, a different instance is online on each of the two servers. If a failure occurs, the instance on the failing node is brought online on the other server, resulting in two instances online on one server. |

Table 8-1 outlines the high-level objectives for implementing the configuration and the tasks for each objective.

Table 8-1 Task list: New High Availability configuration

Objectives	Tasks
“ Reviewing the requirements ” on page 80	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites ■ Review the requirements
“ Reviewing the configuration ” on page 86	Review the configuration
“ Configuring the storage hardware and network ” on page 88	Configure the storage hardware and network
“ Installing Veritas Storage Foundation and High Availability for Windows ” on page 91	<ul style="list-style-type: none"> ■ Verify the driver signing options ■ Install SFW HA ■ Restore driver signing options
“ Configuring disk groups and volumes ” on page 97	<ul style="list-style-type: none"> ■ Planning your storage layout ■ Create disk groups ■ Create volumes ■ Managing disk groups and volumes
“ Configuring the cluster ” on page 106	Use the VCS Cluster Configuration Wizard (VCW) to set up the cluster
“ Installing and configuring the application or server role ” on page 121	<ul style="list-style-type: none"> ■ As necessary, install the application program files on the local drive of the first node ■ Install files relating to the data and logs on the shared storage ■ Deport the disk groups on the first node and import them on the second node ■ Install the application on the second node
“ Configuring the service group ” on page 124	<ul style="list-style-type: none"> ■ Use the applicable wizard to create and configure the VCS service group or groups ■ Bring the service group online
“ Verifying the cluster configuration ” on page 181	<ul style="list-style-type: none"> ■ Switch the service group to the second node ■ Shut down an active cluster node

Table 8-1 Task list: New High Availability configuration (Continued)

Objectives	Tasks
“Possible tasks after completing the configuration” on page 182	<ul style="list-style-type: none">■ Configure the Cluster Management Console connection■ Modify the cluster configuration■ Modify the application or server role service group

Reviewing the requirements

Review these product installation requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

- [“Disk space requirements”](#) on page 80
- [“Requirements for Veritas Storage Foundation High Availability for Windows \(SFW HA\)”](#) on page 80

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 8-2](#) estimates disk space requirements for SFW HA.

Table 8-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 5.1 Service Pack 2 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/358407>
- Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:
<http://entsupport.symantec.com/docs/358406>
- Review the operating systems supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported operating systems

SFW and SFW HA have client and server components that run on specific Windows operating systems.

SFW and SFW HA software for servers

Your system must run one of the operating systems listed in the tables to install the SFW or SFW HA server software:

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 8-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required

Table 8-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

The following table shows supported Windows Server 2008 operating systems.

Note: SFW HA supports Windows Server 2008 R2 without Hyper-V x64: Standard Edition, Enterprise Edition, Datacenter Edition.

Table 8-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server Note: SFW HA does not support Hyper-V.	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2

Table 8-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition Note: SFW HA does not support Hyper-V.	
Windows Server 2008 R2 for IA systems	IA64		
Windows Storage Server 2008			

SFW and SFW HA software for clients

The SFW or SFW HA client software can be installed on any one of the operating system versions, editions, and architectures that the Server Components are supported on, except Server Core.

The following table shows the additional client operating systems that are supported.

Table 8-5 Supported client operating systems

Windows Client	Platform	Edition	Service Pack
Windows XP	x86, x64		SP3 required
Windows Vista	x86, x64,	Ultimate Edition, Business Edition, Premium Edition	SP1 or SP2 required
Windows 7	x86, x64	Ultimate Edition, Business Edition, Premium Edition	

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs. See "[Best practices for SFW HA](#)" on page 86.
- NIC teaming is not supported for the VCS private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each application virtual server.
 - A minimum of one static IP address for each physical node in the cluster.

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements for SFW HA

The following permissions are required:

- You must be a domain user.
- You must be a member of the local Administrators group on all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements for SFW HA

Please review the following additional requirements:

- Installation media for all products and third-party applications.
- Licenses for all products and third-party applications.

- You must install the operating system in the same path on all systems. For example, if you install Windows Server on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- For a Replicated Data Cluster, install only in a single domain.

Best practices for SFW HA

Symantec recommends that you perform the following tasks:

- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the vxclus UseSystemBus ON command. This is applicable for a Replicated Data Cluster configuration.

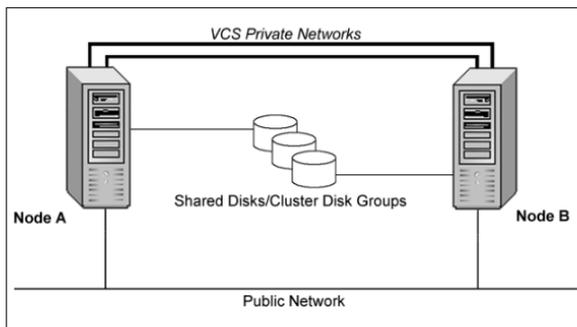
Reviewing the configuration

This example configuration is one of the most common configurations for a cluster. It is a new installation with two servers and one storage array, in an Active-Passive configuration where the active node of the cluster hosts the virtual server and the second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. The example describes a generic database application.

The example configuration does not include dynamic multi-pathing.

See "[Adding DMP to a clustering configuration](#)" on page 67.

Figure 8-1 SFW HA Active-Passive configuration with two servers



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.

Note: Follow the appropriate procedure for verifying the DNS settings, as the procedures for Windows Server 2003 and Windows Server 2008 are slightly different.

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 VCS 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.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 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.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 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 and Dial-up 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 VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in 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**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 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 VCS 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.
- 13 Click **OK**.

Installing Veritas Storage Foundation and High Availability for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 8-6 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 92.

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.

Installing Symantec Trusted certificate for unsigned drivers

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 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.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

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 HA, 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, select the appropriate SFW product options, and click **Next**.

Client Components	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used while configuring high availability.
-------------------	--

	Required to install the Solutions Configuration Center which provides information and wizards to assist configuration
--	---

Global Cluster Option	Required for a disaster recovery configuration only.
-----------------------	--

Volume Replicator (VVR)	If you plan to use Veritas Volume Replicator (VVR) for replication, select this option to install VVR.
-------------------------	--

High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.
---	---

- 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.

If you are using multiple paths and selected 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 path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation for Windows (Server Components) with the Microsoft Cluster Server (MSCS) option, 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. For more information, see the *Veritas Storage Foundation Administrator's Guide*.

Review the messages, and then click **OK**.

- 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.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

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 disk groups and volumes

Use Veritas Storage Foundation for Windows to create cluster disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 97
- [“Creating dynamic cluster disk groups”](#) on page 99
- [“Creating dynamic volumes”](#) on page 101

Planning disk groups and volumes

The requirements for disk groups and volumes depend on the type of application or server role. Review the requirements and best practices for your application or server role:

- [Planning your File Share storage](#)
- [Planning your IIS storage](#)
- [Planning your Microsoft Virtual Machine storage](#)
- [Planning your storage for additional applications](#)

Planning your File Share storage

Considerations for planning the File Share storage:

- Make sure that the disk group and volumes for the file server shared directory are configured on shared storage.
- When configuring a new set up, first create the disk groups and volumes on the shared storage and then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.

Planning your IIS storage

Considerations for planning the IIS storage:

- Make sure that the disk groups and volumes which will host the directory and files for the web sites are on the shared storage.
- For a new IIS installation, make sure that the directory for the web sites is created on volumes on the shared storage.
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the web site in IIS and then restart the web site again.

Planning your Microsoft Virtual Machine storage

Make sure the volumes that contain the shared virtual disk files for the virtual machines are located on the shared storage.

Planning your storage for additional applications

The information provided in this section is generic to any application. Make sure you create the appropriate disk groups and volumes to hold the application data. If your application requires replication of registry keys between the cluster systems, then Symantec recommends that you create a dedicated RegRep volume so that its MountV dependency is not linked with any other application-specific resources in the group.

Decide how you want to organize the disk groups and the number and type of volumes you want to create. Some considerations are:

- The number of disk groups that are needed
The number of disk groups depends on your application and the planned organization of the data. VCS requires that the application program files be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application would be contained in a single disk group.
- The type of volumes you want to create
 - Mirrored and RAID-5 volumes provide fault tolerance for critical data.
 - Striped volumes add performance capability.
 - Volumes that are both mirrored and striped offer both performance and fault tolerance.

Note: If you plan to use replication software, such as VVR, do not use software RAID-5 volumes. This does not apply to hardware RAID-5.

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.
- The implications of backup and restore operations for the disk group setup.
- The sizes of databases and logs, which depend on the traffic load.
- If you plan to implement a disaster recovery configuration with VVR, a Storage Replicator Log (SRL) volume is required for each disk group that contains volumes that are replicated. You can create the SRL volume now or you can create it later when you run the Disaster Recovery Wizard. If you create it later, ensure that you allow sufficient disk space for this volume.

Creating dynamic cluster disk groups

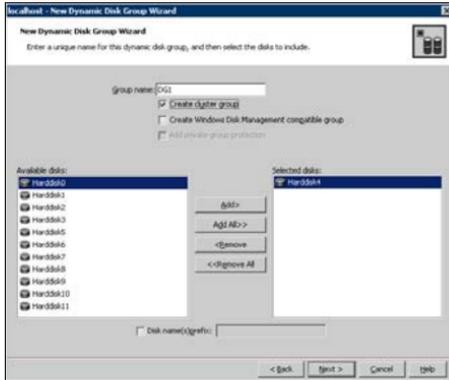
Follow the steps in this section to create one or more disk groups for your application.

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 disk group name (for example, DG1).
- 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.

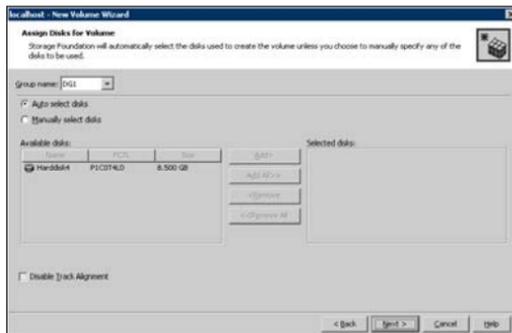
- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
- 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

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.
- 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.

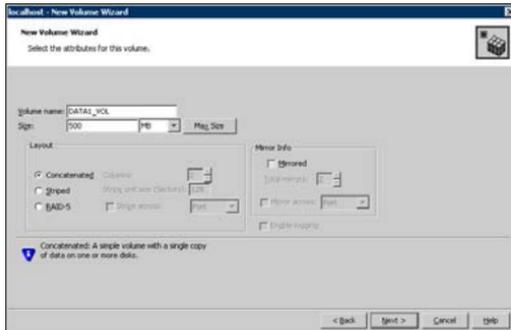


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

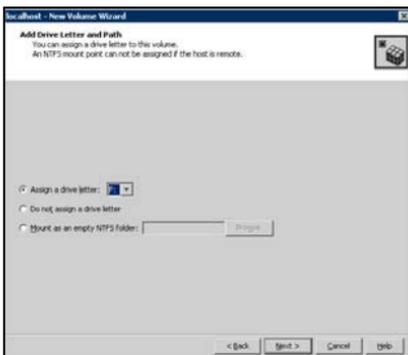
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the volume attributes.

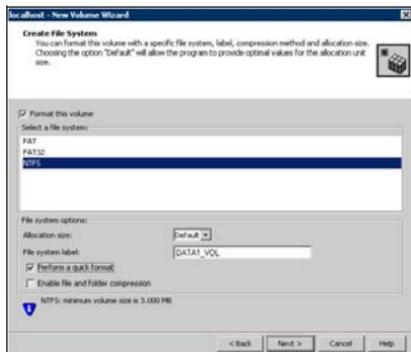


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** 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 additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Managing disk groups and volumes

This section includes the following procedures:

- Importing a disk group and mounting a shared volume
- Unmounting a volume and deporting a disk group

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**.

Configuring the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

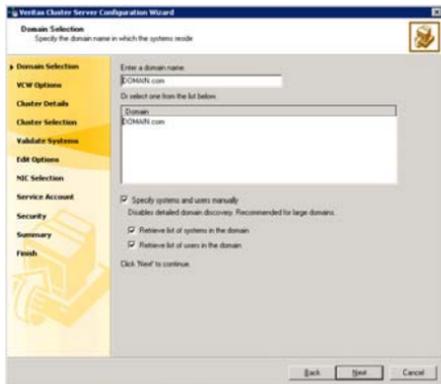
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.

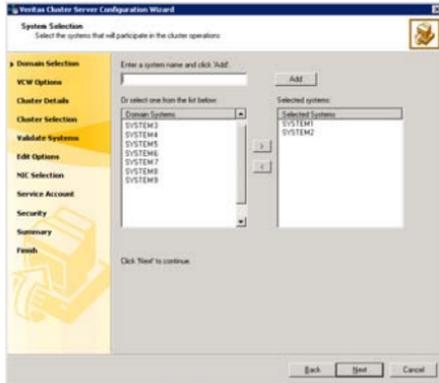
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 108.
 - To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**. If you chose to retrieve the list of systems, proceed to [step 6](#) on page 108. Otherwise, proceed to the next step.
- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster. Proceed to [step 8](#) on page 108.

- 6 On the System Selection panel, specify the systems for the cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

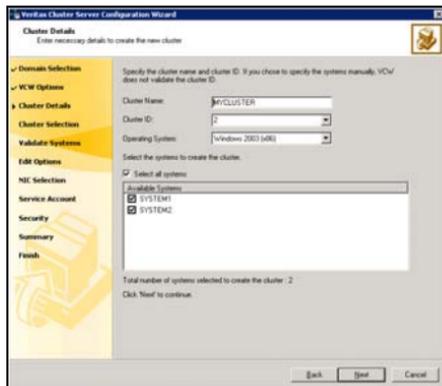
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

Operating System From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (NICs) on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

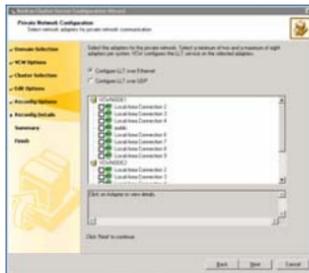
- 10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 112.

- 11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



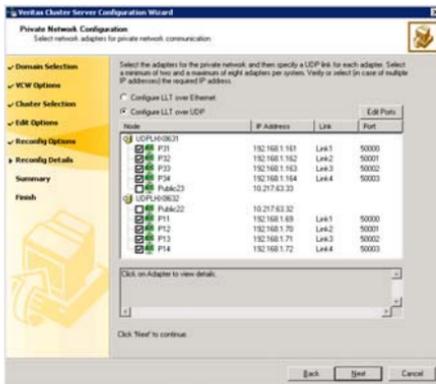
- Select the check boxes next to the two NICs to be assigned to the private network.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links.

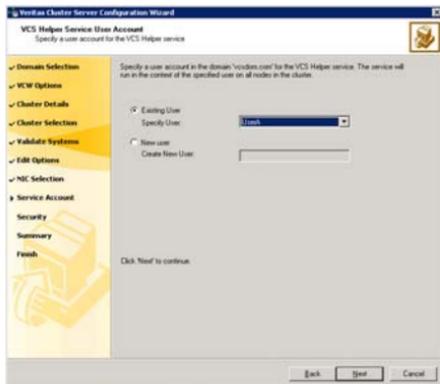
Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

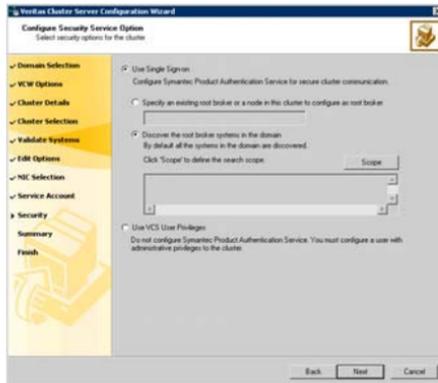
- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 107, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:
- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers.

Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user

Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 8-7 contains some more examples of search criteria.

Table 8-7 Search criteria examples

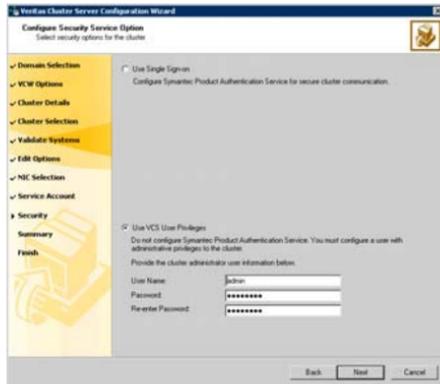
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
 - Click **Next**.
- 14 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

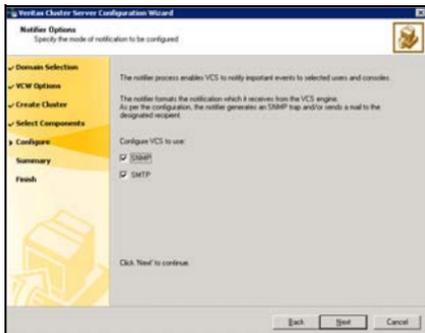
- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 117.

Configuring notification

This section describes steps to configure notification.

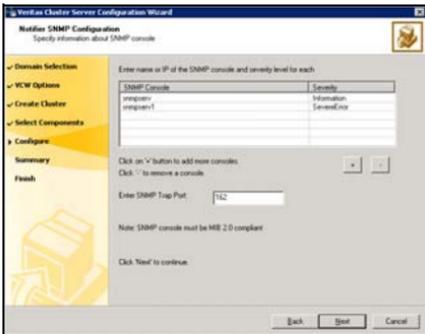
To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

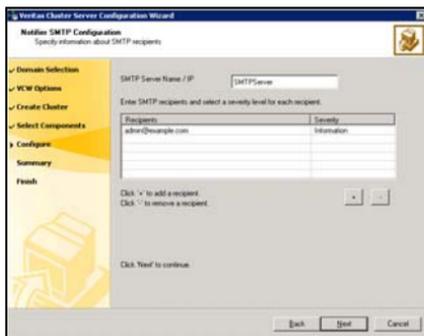


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

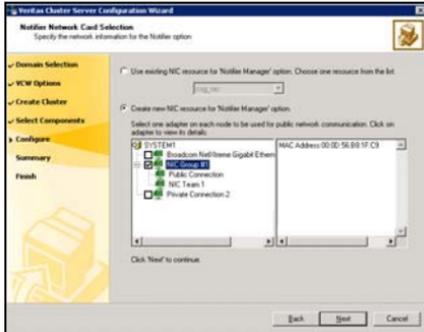


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

Installing and configuring the application or server role

This section provides considerations for installing and configuring your application or server role.

See the following topics:

- [Configuring a File Share server role](#)
- [Configuring a Print Share server role](#)
- [Installing and configuring the IIS application](#)
- [Installing and configuring Microsoft Virtual Server](#)
- [Installing additional applications](#)

Configuring a File Share server role

Points to note when configuring a File Share:

- Make sure that the disk group and volumes for the file server shared directory are configured on shared storage.
- When configuring a new set up, first create the disk groups and volumes on the shared storage and then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.
- The FileShare agent is installed automatically with SFW HA.

Configuring a Print Share server role

Points to note when configuring a Print Share:

- Make sure the printer is connected to the network and is configured with an IP address.
- Install software drivers for the network printer on all systems in the cluster.

To add a print driver

- 1 Open the **Printers** Control Panel.
- 2 Click **File > Server Properties**.
- 3 In the Print Server Properties dialog box, click the **Drivers** tab.

- 4 Click **Add**. This launches the Add Printer Driver wizard.
- 5 Follow the wizard instructions to add the printer driver on the system. You must add the driver on each system that will be part of the service group.

Installing and configuring the IIS application

Points to note when installing IIS:

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on shared storage.
- Import the cluster disk groups and mount the volumes that contain the website data, on the first node.
- For a new IIS installation, while creating new web sites, create the site folder on the shared storage and place the site content in that folder.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the website in IIS and then restart the website again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group.

To synchronize the IIS configuration on Windows 2003 systems

- 1 Synchronize the IIS configuration on all nodes that will host the IIS service group. Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system.

For example, the following command copies the IIS metabase to *target_system*. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing and configuring Microsoft Virtual Server

Points to note when installing MS Virtual Server:

- Verify Microsoft Virtual Server is installed and configured identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- Install and configure Virtual Machine Additions *on each virtual machine* if you plan to enable detailed monitoring for the virtual machine resources.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

Installing additional applications

Following are some very generic points for installing any application:

- Make sure that the disk groups and volumes are mounted on the node before installing the application.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes must be on drive C.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the shared storage.

Configuring the service group

The Solutions Configuration Center provides wizards to configure the service groups for the additional SFW HA applications or server roles. It also supports the Application Configuration Wizard which can be used to configure any other application for which application specific wizards have not been provided. Depending on the application that you have installed, complete the appropriate procedure to configure the service group:

- [Configuring the File Share service group](#)
- [Configuring the PrintShare service group](#)
- [Configuring the IIS service group](#)
- [Configuring the MSVirtual Machine service group](#)
- [Configuring the service group for any additional applications](#)
- [Configuring an Oracle service group](#)

Configuring the File Share service group

Configuring the File Share service group involves creating a FileShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically
`C:\Program Files\Veritas\Cluster Server.`
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
 - Verify that the directories to be shared reside on shared drives.

- Mount the drives containing the shared directories on the system where you run the wizard. Unmount the drives from other systems in the cluster.
- Verify that Veritas Command Server service is running on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - A unique virtual computer name to be assigned to the file share server. This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
 - A unique virtual IP address to be assigned to the file share server. This is the IP address that the clients can use to access the file share.

Note: Windows Server 2008 does not support accessing file shares using the virtual IP address.

- The list of directories to be shared.

The wizard enables you to add existing shares to the VCS configuration. However, you cannot add special shares (shares created by the operating system for administrative and system use). For example, you cannot add the shares ADMIN\$, print\$, IPC\$, and *DriveLetter\$* to the VCS configuration.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

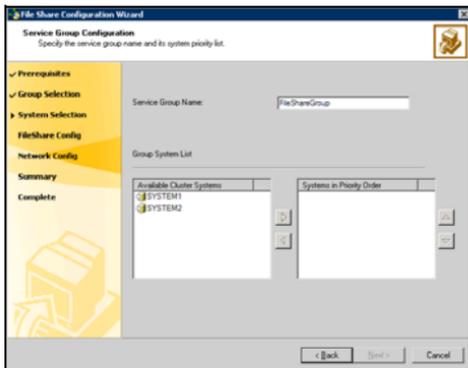
- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).

Creating the FileShare service group

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide*, for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

To configure a FileShare

- 1 Start the File Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > File Share Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > File Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the File Share service group.

Available Cluster Systems

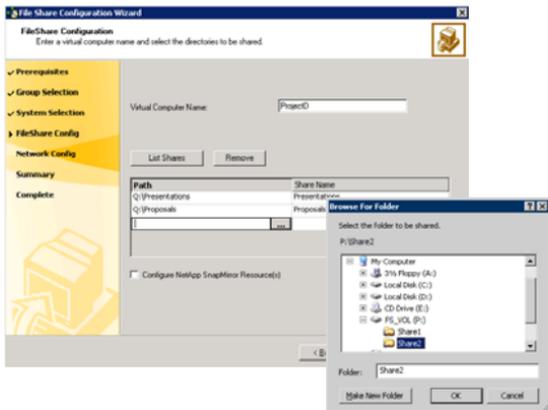
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

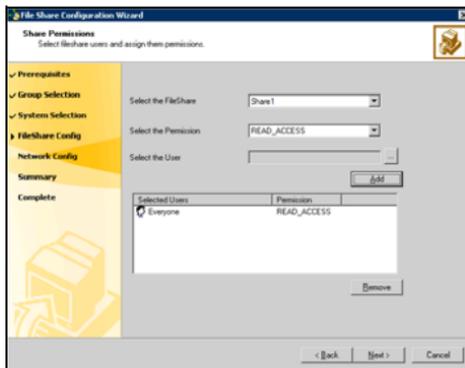
- 5 On the File Share Configuration panel, specify the configuration information for the FileShare resources to be created and then click **Next**. The wizard begins validating your configuration. Various messages indicate the validation status.

**Virtual Computer Name**

Type a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 16 characters.

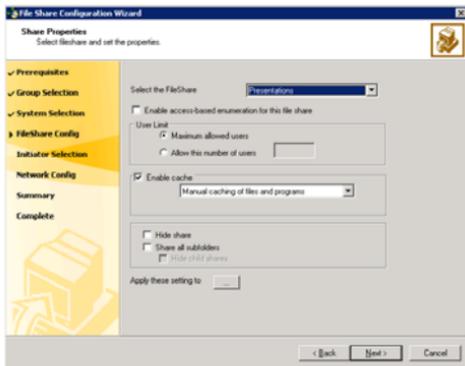
List Shares	Click List Shares to view the existing shares on the shared storage, then select a share and click Add . You cannot add special shares (shares created by the operating system for administrative and system use).
Path	Click the field and either type the path of the directory to be shared or click the ellipsis button (...) to browse for a directory. The selected directories must meet the following conditions: <ul style="list-style-type: none">■ The selected drive, the mount path, and the file path must not exist in the VCS configuration.■ The directories to be shared must reside on shared, non-system drives. The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.
Share Name	If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can select a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.
Remove	To remove a file share from the configuration, click to select the file share and then click Remove .

- 6 On the Share Permissions panel, select the file share, specify the users for the file shares and assign permissions to them, and then click **Next**.



- | | |
|-----------------------|--|
| Select the FileShare | From the drop-down list, select the file share with which to associate user permissions, or select the default All FileShares to set the same permissions for all file shares. |
| Select the Permission | From the drop-down list, select the permission to be associated with the user. |
| Select the User | Click the ellipsis button (...), select a user, and click OK . |
| Add | Click Add to add the specified user to the Selected Users list. By default, all selected users are given READ_ACCESS permission. |
| Selected Users | Displays a list of selected users and the file share permissions. You can configure a maximum of 50 users for each file share. To configure more users, create a user group.
To change the file share permission associated with a user, click a user name in the Selected Users list and then select the desired permission from the Select the Permission drop-down list. |
| Remove | To deny file share access to a user, click the user name in the Selected Users list and click Remove . |

- 7 On the Share Properties panel, set the share properties for the file shares and then click **Next**.



Select the FileShare

From the drop-down list select a file share whose properties you wish to set.

Enable access-based enumeration for this file share

Check the **Enable access-based enumeration** check box to enable the Windows access-based enumeration feature on the selected file share.

User Limit

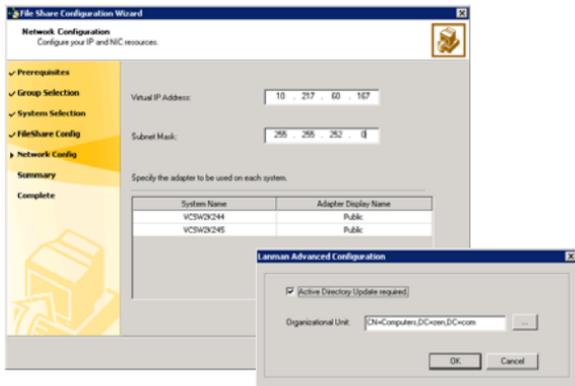
Specify the number of users that are allowed access to the selected file share.

Choose from the following options:

- **Maximum allowed users:** Select this option to allow access to the maximum numbers of users allowed on Windows.
- **Allow this number of users:** Select this option and then type the number of users that you wish to grant access to the selected file share.
If you type zero or a value greater than what Windows supports, access is granted to the maximum allowed users on Windows.

Enable cache	<p>Check the Enable cache check box to enable local caching of the contents of the selected file share. Then, specify how the contents of the file share are available to users for offline access. In the drop down list select from the following caching options:</p> <ul style="list-style-type: none">■ Manual caching of files and programs: Only the files and programs specified by the user are available offline. This sets the FileShare resource attribute ClientCacheType to MANUAL.■ Automatic caching of programs: All the files and programs that the users access from the file share are available offline. This sets the FileShare resource attribute ClientCacheType to DOCS.■ Optimized automatic caching of files and programs: All files and programs, including executables, are cached locally. The next time the user accesses the executable files, they are launched from the local cache. This sets the FileShare resource attribute ClientCacheType to PROGRAMS.
Hide share	Check the Hide Share check box to make the new share a hidden share.
Share all subfolder	Check the Share all subfolders check box to share the subdirectories.
Hide child shares	Check the Hide child shares check box to hide the shared subdirectories.
Apply these settings to	<p>To apply the specified share properties to multiple file shares simultaneously, do the following:</p> <ol style="list-style-type: none">1 Click the ellipsis (...) button.2 On the Copy Share Properties dialog box, select the file shares from the Available Shares list and click the right arrow to move them to the Selected Shares list. Note that only those files shares that are not already shared are available for selection.3 Click OK. <p>Note: This option is not visible if you are configuring only one share in the service group.</p>

- On the Network Configuration panel, specify information related to your network and then click **Next**.



Virtual IP Address

Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual IP belongs.

Adapter Display
Name

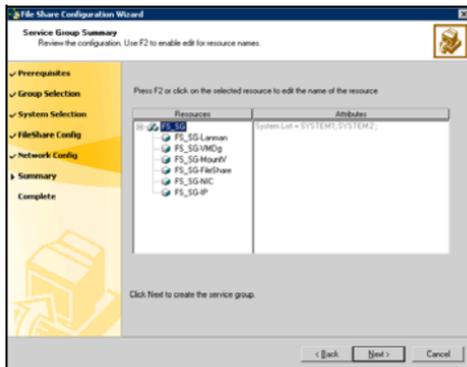
For each system in the cluster, select the public network adapter name.

This field displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. Verify that you select the adapters assigned to the public network, not the private.

Advanced Settings Click **Advanced Settings** to specify additional details for the Lanman resource.
On the Lanman Advanced Configuration dialog box, complete the following:

- 1 Check **Active Directory Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
 - 2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format
CN=containername,DC=domainname,DC=com.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
By default, the Lanman resource adds the virtual server to the default container "Computers."
 - 3 Click **OK**.
The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- 9 On the Service Group Summary panel, review the service group configuration and click **Next**.
Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

- 10 In the completion dialog box, check **Bring the service group online** check box if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the PrintShare service group

Configuring the Print Share service group involves creating a PrintShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator* mode and then run the VCS commands.

Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Verify that VCS Command Server is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.
- Symantec recommends creating spooler and the replication directories on different disk partitions or volumes.
- Mount the drives with the spooler and the replication directories on the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).

- Verify that the software drivers for the network printers are installed on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - A unique virtual computer name to be assigned to the print share server.
This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
 - A unique virtual IP address to be assigned to the print share server.
 - The network printer's IP address.

Creating the PrintShare service group

To create a Print Share service group perform the following tasks:

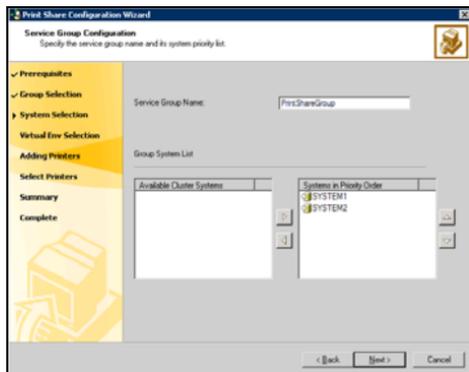
- Create a new service group with a PrintSpool resource and bring it online. This also involves configuring the Lanman resource on which the PrintSpool resource depends.
- Add a network printer to the virtual computer created by the Lanman resource. Create a new TCP/IP port for the printer.
- Configure a PrintShare resource in your service group and bring it online.

The following procedures describe these tasks in detail.

To create a Print Share service group with a PrintSpool resource

- 1 Start the Print Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Print Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

- 4 On the Service Group Configuration panel, specify the service group details and click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the Print Share service group.

Available Cluster Systems

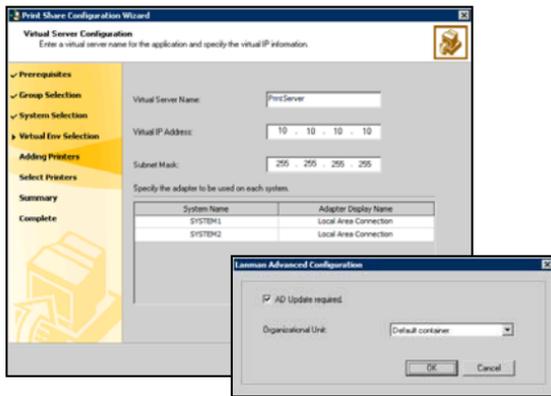
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Virtual Server Configuration panel, specify information related to your network and then click **Next**.



Virtual Server Name

Type a unique virtual computer name by which the server will be known to clients. Note that the virtual name must not exceed 16 characters.

Virtual IP Address

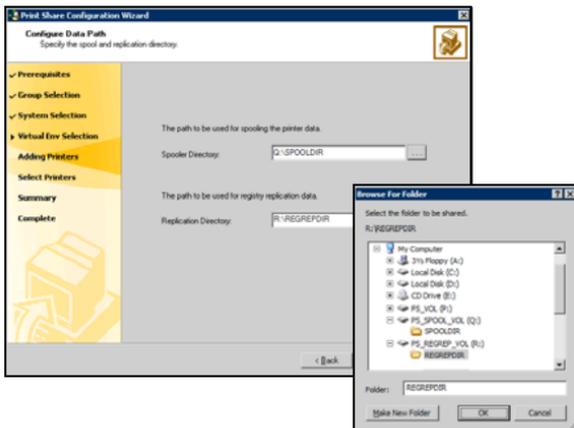
Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual server belongs.

Advanced Settings	<p>Click Advanced Settings to specify additional details for the Lanman resource.</p> <p>On the Lanman Advanced Configuration dialog box, complete the following:</p> <ol style="list-style-type: none"> 1 Check AD Update required check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true. 2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format CN=containername,DC=domainname,DC=com. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers." 3 Click OK. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
Adapter Display Name	<p>Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow.</p> <p>For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.</p>

- On the Configure Data Path panel, specify the spool and registry replication directories and then click **Next**.



Spooler Directory

Type the path or click ... (ellipsis button) to browse for the directory. All print commands will be spooled at this location.

Replication Directory

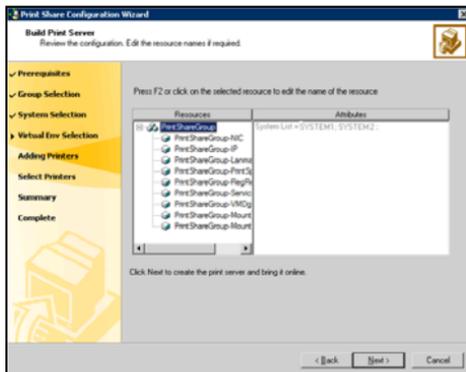
Type the path or click ... (ellipsis button) to browse for the directory. All changes related to the printer registry keys will be logged at this location.

The selected directories must fulfill the following conditions:

- The selected drive, the mount path, and the file path must not exist in the VCS configuration.
- The directories to be shared must reside on shared, non-system drives.
Symantec recommends creating the directories for replication and spooling on different mounts.

- On the Build Print Server panel, review the configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**. The wizard starts running commands to add the PrintSpool resource and the resources on which the

PrintSpool resource depends, including the Lanman and ServiceMonitor resources.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
- To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

8 Bring the PrintSpool resource online.

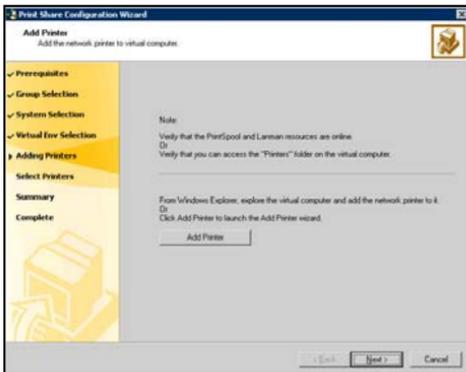
Proceed to the next step to add the network printer to the virtual computer created by the Lanman resource and to create a new TCP/IP port for the printer.

To add the network printer to the virtual computer

- 1 Launch the Add Printer wizard to add the network printer to the virtual computer. Before starting the Add Printer wizard, verify that the PrintSpool and Lanman resources are online in your configuration.

To launch the Add Printer wizard, return to the Print Share Configuration Wizard and click **Add Printer** on the Add Printer panel, or in Windows Explorer, search for the virtual computer, explore the virtual computer by

double-clicking its name and on the virtual computer's Printers folder, double-click **Add Printer**.

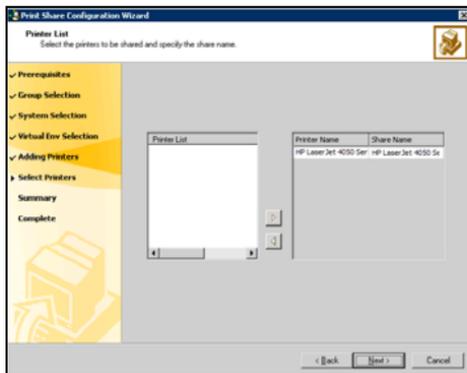


- 2 In the Add Printer wizard, review the information in the Welcome panel and click **Next**.
- 3 Follow the wizard instructions to add the network printer to the virtual computer.
In the Printer Sharing dialog box, always choose the **Do not share this printer** option.
Repeat these steps for each additional printer to be installed.
- 4 Return to the Print Share Configuration Wizard, and proceed to the next step to configure a PrintShare resource in your service group and bring it online.

To configure a PrintShare resource for the service group

- 1 On the Add Printer panel, click **Next**.

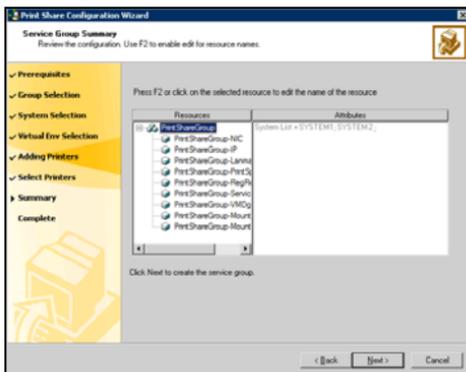
- On the Printer List panel, specify the printers to be included in the Print Share service group and then click **Next**.



- Printer List** Click to select the printer, and then click the right arrow to include the selected printers in your service group. To remove a selected printer from your service group, click the printer from the Printer Name list and click the left arrow.
- Share Name** Type a unique share name for the printer by which it will be known to clients. If you previously chose to share the printer, VCS uses the printer's share name.

- On the Service Group Summary panel, review the service group configuration and then click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration.

Click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



Resources Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes Displays the attributes and their configured values, for a resource selected in the Resources list.

- 4 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the IIS service group

Configuring the IIS service group involves creating a IIS service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify the sites to be monitored are on shared storage.
- For IIS 7.0 on Windows Server 2008, you must install the following components:
 - IIS 6 Metabase Compatibility
 - IIS 6 WMI CompatibilityThese options are available under Management Tools on the Role Services page of the Add Roles Wizard. These components are required for the IIS agent to function on Windows Server 2008.
For IIS 7.0 on Windows Server 2008 Server Core, you must install IIS in the specified order. See [“Installing IIS 7.0 on Windows Server 2008 Server Core”](#) on page 147 for instructions.
- Do not use the IIS agent to configure SMTP and NNTP sites if you have Microsoft Exchange installed.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- Verify the port numbers assigned to IIS sites are not used by other programs.
- Synchronize the IIS configuration on all nodes hosting the service group. See [“Synchronizing IIS configuration on Windows 2003”](#) on page 146 for instructions.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically C:\Program Files\Veritas\Cluster Server.
 - Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- To configure IIS agent on Windows Server 2008 Server Core, you must manually add the required resources and configure the service group. You can perform the manual configuration steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console).

Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Mount the drives containing the shared directories from the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - IIS sites to be monitored.
 - Application pools associated with each site.
 - Port numbers associated with each site.
 - Virtual IP addresses and computer names associated with the sites. The virtual IP addresses and the virtual computer names must have forward and reverse entries in the DNS.
- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Synchronizing IIS configuration on Windows 2003

Complete the following steps.

To synchronize the IIS configuration on Windows 2003 systems

Synchronize the IIS configuration on all nodes that will host the IIS service group.

- 1 Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system. For example, the following command copies the IIS metabase to `target_system`. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing IIS 7.0 on Windows Server 2008 Server Core

Complete the following steps.

To install IIS 7.0 on Windows Server 2008 Server Core

- 1 Type the following at the command prompt:

```
start /w pkgmgr  
/iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;  
IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;  
IIS-HttpErrors;IIS-HttpRedirect;IIS-ApplicationDevelopment;  
IIS-ASP;IIS-CGI;IIS-ISAPIExtensions;IIS-ISAPIFilter;  
IIS-ServerSideIncludes;IIS-HealthAndDiagnostics;  
IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;  
IIS-HttpTracing;IIS-CustomLogging;IIS-ODBCLogging;IIS-Security;IIS-BasicAuthentication;IIS-WindowsAuthentication;  
IIS-DigestAuthentication;  
IIS-ClientCertificateMappingAuthentication;  
IIS-IISCertificateMappingAuthentication;  
IIS-URLAuthorization;IIS-RequestFiltering;IIS-IPSecurity;  
IIS-Performance;IIS-HttpCompressionStatic;  
IIS-HttpCompressionDynamic;IIS-WebServerManagementTools;  
IIS-ManagementScriptingTools;IIS-IIS6ManagementCompat
```

```
ibility;  
IIS-Metabase; IIS-WMICompatibility; IIS-LegacyScripts;  
IIS-FTPPublishingService; WAS-WindowsActivationService  
;  
IIS-FTPPublishingService; IIS-FTPServer
```

- 2 Verify that all the components specified in the earlier step have successfully installed. Type the following at the command prompt:
`notepad C:\windows\logs\cbs\cbd.log`
This opens the log file, `cbd.log`, in the Notepad text editor.
- 3 Check the entries in the log file, `cbd.log`. The last log entry should resemble the following:
`Info CBS Pkgmgr: return code: 0x0`
This message indicates that all the components are installed successfully.
- 4 Run the `oclist` command to verify that the following components are installed:
`IIS-WebServerRole; IIS-WebServer; IIS-IIS6ManagementCompatibility;
IIS-Metabase; IIS-WMICompatibility; IIS-FTPPublishingService;
WAS-WindowsActivationService; IIS-FTPPublishingService; IIS-FTPServer`
Type the following at the command prompt:
`oclist`
- 5 Repeat the steps on all the nodes on which you wish to configure the IIS service group.

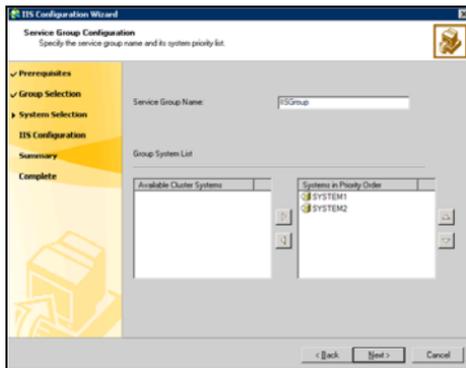
Creating the IIS service group

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

To create an IIS service group

- 1 Start the IIS Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > IIS Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > IIS Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.

- 3 On the Wizard Options panel, click **Create service group** and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the IIS service group.

Available Cluster Systems

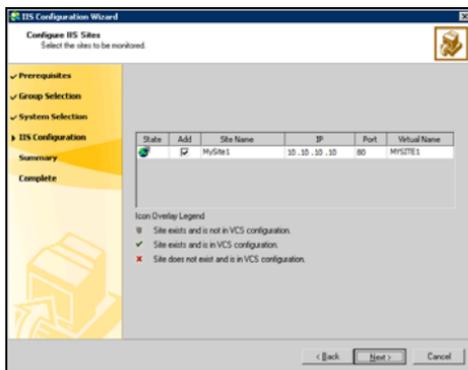
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system in the Systems in Priority Order and click the up and down arrows.

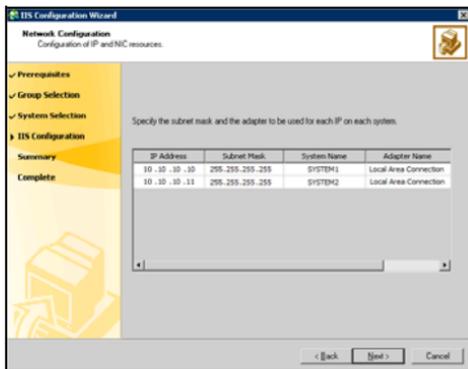
System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names, and then click **Next**.



- | | |
|------|--|
| Add | Check the check box corresponding to the site to be configured in VCS. |
| IP | Type the virtual IP address for each site to be configured. Make sure that each virtual IP address is associated with only one virtual computer name and vice-versa. |
| Port | Type the port number for each site to be configured. |

- 6 On the Network Configuration panel, specify information related to the virtual IP addresses and click **Next**.



IP Address

Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems.

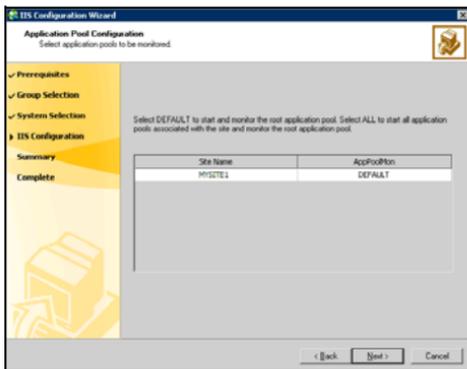
Subnet Mask

Type the subnet mask associated with each virtual IP address.

Adapter Name

Select the adapter associated with the virtual IP address on each system.

- 7 On the Application Pool Configuration panel, select the monitoring options for application pools associated with each site and click **Next**.

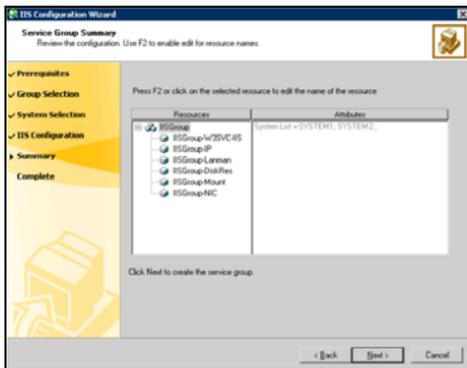


Site Name Displays the site names.

AppPoolMon For each site, select the monitoring options from the AppPoolMon list.
NONE—The agent will not monitor the application pool associated with the site.
DEFAULT—Starts and monitors the root application pool associated with the site.
ALL—Starts all application pools associated with the site and monitors root application pool.

- 8 On the Service Group Summary panel, review the service group configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click

Yes. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
- To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

- In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the MSVirtual Machine service group

Configuring the MSVirtual Machine service group involves creating a MSVirtual Machine service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- Verify that the shared drives required by the applications are mounted.
- Mount the drives containing the shared directories of the virtual machine, on which the wizard will be run. Unmount the drives from other systems in the cluster.

- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
- Disable the firewall on each node that will host the service group.
- You must have the following information ready. The wizard will prompt you for this information:
 - The name of the virtual machine.
 - Destination on shared disks for the virtual hard disk files.
 - Network adapters on physical nodes to be associated with network adapters on the virtual machine.
 - Information about monitoring heartbeats (optional).

Creating the MSVirtual Machine service group

The following procedure describes how to create a MSVirtual Machine service group using the MSVirtual Machine Configuration Wizard.

To create the MSVirtualMachine service group

- 1 Start the MSVirtual Machine Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > MSVirtual Machine Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > MSVirtual Machine Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, select the **Create service group** option and click **Next**.
- 4 Enter a name for the service group and specify the systems on which to configure the service group.
 - Enter a name for the service group.
 - In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the **Systems in Priority Order** box and click the left arrow.

- To change a system's priority in the service group's system list, click the system from the **Systems in Priority Order** and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
- Click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.

5 Specify details about the virtual machine.

- Select the virtual machine.
- For each virtual disk, specify a destination folder where the virtual hard disk files will be moved. Click the Browse icon to browse for folders.
- To enable detail monitoring for the virtual machine, select the **Monitor Heartbeats** check box and enter failed heartbeat threshold in the **No. of Monitor Cycles** field.

The threshold defines the number of consecutive monitor cycles the agent waits to detect heartbeats from the virtual machine before declaring the resource as faulted.

- Click **Next**.

6 Select an adapters corresponding to the virtual machine on each system.

- For each system in the cluster, enter or click a network adapter name to be associated with the network adapters on the virtual machine. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

The fields for the virtual IP address and subnet mask are disabled by design.

- Click **Next**.

7 Review the service group configuration.

The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.

- The wizard assigns unique names to resources. Change names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.

- Click **Next**.

- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.
The wizard starts running commands to create the service group. Various messages indicate the status of these commands.
- 8 In the completion dialog box, select the check box if you want to bring the service group online on the local system.
- 9 Click **Finish**.

Configuring the service group for any additional applications

Configuring the service group for any additional application involves creating an application service group and defining the attribute values for its resources. This can be done using the Application Configuration Wizard. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- The application is installed on all the nodes that are going to be part of the service group.
 - The shared drives required for the application are mounted on this node.
 - The startup type of the application service is set to manual on all nodes in the service group.
 - The application service is stopped on all nodes in the service group.
 - If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the

steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console).

Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Before running the wizard, make sure you have the following information ready:
 - Details (service name, start parameters, startup directory, etc.) of the application that you wish to configure.
 - Shared storage used by the applications.
 - Application registry entries for configuring registry replication.
 - Network and virtual computer (Lanman) details for the application.

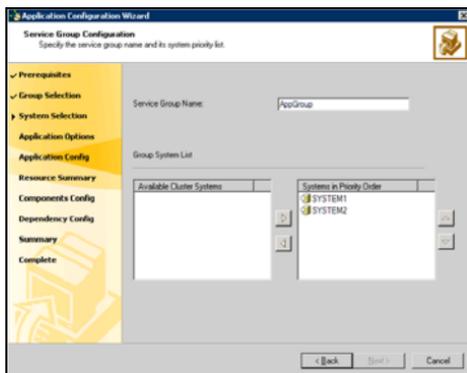
Creating the application service group

The following procedure describes how to create an application service group using the Application Configuration Wizard.

To create an application service group

- 1 Start the Application Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Application Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Application Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Create service group** and click **Next**.

4 Specify the service group name and system list.

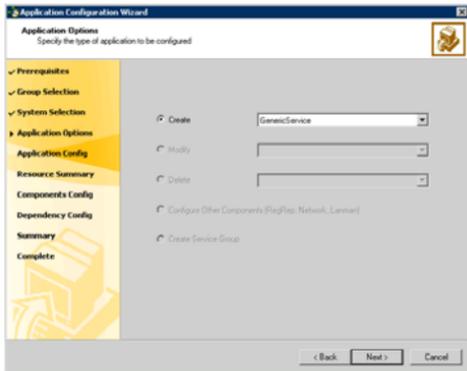


- Enter a name for the service group.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.
To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
 - To change a system's priority in the service group's system list, select the system in the Systems in Priority Order list and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 The Application Options panel provides you the option to specify the type of application to be configured. The available options are:
- **Generic Service:** Configures a service using the Generic Service agent. The agent brings services online, takes them offline, and monitors their status. See "[Configuring a GenericService resource](#)" on page 160.
 - **Process:** Configures a process using the Process agent. The agent brings processes online, takes them offline, and monitors their status. See "[Configuring processes](#)" on page 164.
 - **Service Monitor:** Configures a service using the ServiceMonitor agent. The agent monitors a service or starts a user-defined script and

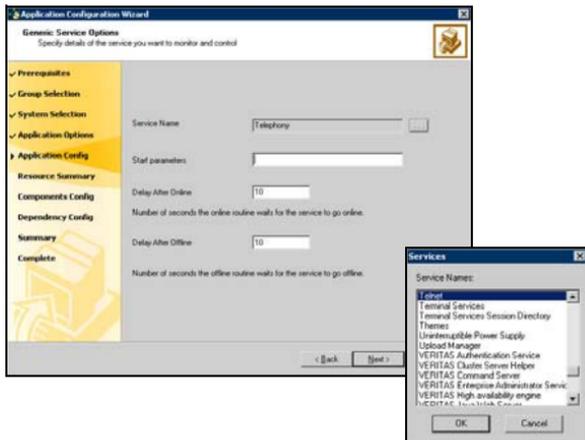
interprets the exit code of the script. See “[Configuring a ServiceMonitor resource](#)” on page 168.

Configuring a GenericService resource

- 1 In the Application Options panel, click **Create**, select **GenericService** from the corresponding drop-down list, and click **Next**.

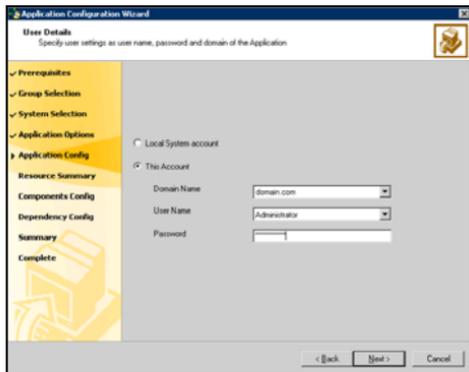


- 2 Select the service name for which you wish to configure a GenericService resource. Also specify the attributes for the resource.



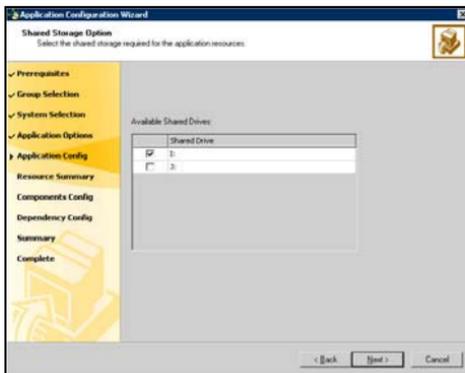
- Click the icon (...) adjacent to the Service Name text box.
- In the Services dialog box, select a service and click **OK**. The selected service appears in the Service Name text box.
- In the **Start Parameters** text box, provide the start parameters for the service, if any.
- In the **Delay After Online** text box, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
- In the **Delay After Offline** text box, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
- Click **Next**.

- 3 Specify the information about the user in whose context the service will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the GenericService resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

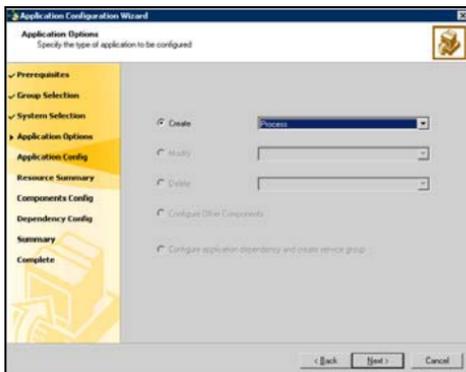


- In the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another GenericService resource, repeat [step 1](#) through [step 5](#).
 - To configure a Process resource, proceed to “[Configuring processes](#)” on page 164 for instructions.
 - To configure a ServiceMonitor resource, proceed to “[Configuring a ServiceMonitor resource](#)” on page 168 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 171 for instructions.

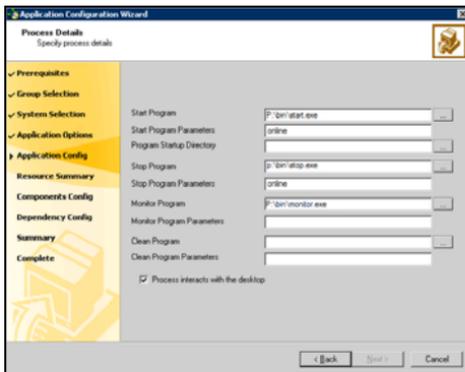
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 174.

Configuring processes

- 1 In the Application Options panel, click **Create**, select **Process** from the corresponding list, and click **Next**.



- 2 Specify the details for the process.



- In the **Start Program** text box, specify the complete path of the program that will start the process to be monitored by VCS. You can

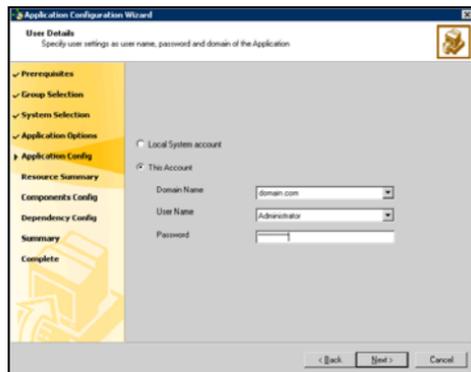
choose to either type in the location of the program or browse for it using the (...) icon.

- In the **Start Program Parameters** text box, specify the parameters used by the Process agent start program.
- In the **Program Startup Directory** text box, enter the complete path of the Process agent program or browse for it by clicking the (...) icon.
- In the **Stop Program** text box, enter the complete path of the program that will stop the process started by the Start Program or browse for it by clicking the (...) icon.
- In the **Stop Program Parameters** text box, specify the parameters used by the stop program.
- In the **Monitor Program** text box, enter the complete path of the program that monitors the Start Program or browse for it by clicking the (...) icon.

If you do not specify a value for this attribute, VCS monitors the Start Program. If the Start Program is a script to launch another program, you must specify a monitor program.

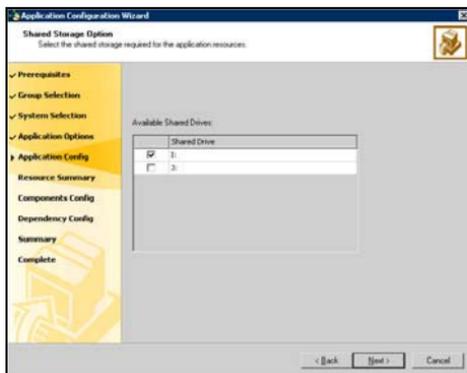
- In the **Monitor Program Parameters** text box, specify the parameters used by the monitor program.
- In the **Clean Program** text box, enter the complete path of the Clean process or browse for it by clicking the (...) icon.
- If no value is specified, the agent kills the process indicated by the Start Program.
- In the **Clean Program Parameters** text box, specify the parameters used by the Clean program.
- Select the **Process interacts with the desktop** check box if you want the process to interact with your Windows desktop. Setting this option enables user intervention for the process.
- Click **Next**.

3 Specify information about the user in whose context the process will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the Process resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

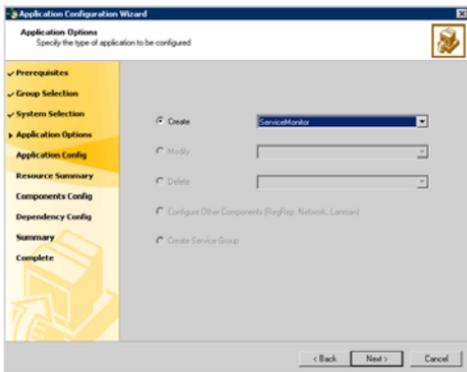


- From the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the Process resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another Process resource, repeat [step 1](#) through [step 5](#).
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 160 for instructions.
 - To configure a ServiceMonitor resource, proceed to [“Configuring a ServiceMonitor resource”](#) on page 168 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to [“Configuring VCS components”](#) on page 171 for instructions.

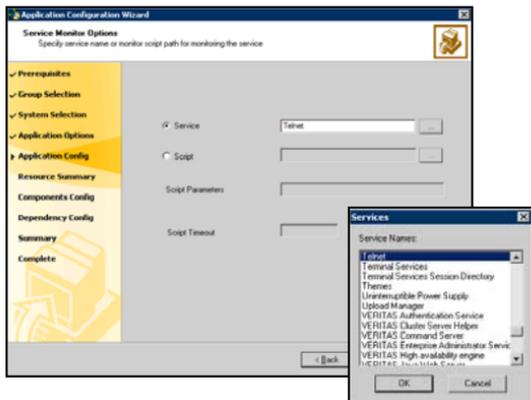
If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 174.

Configuring a ServiceMonitor resource

- 1 In the Application Options panel, click **Create**, select **ServiceMonitor** from the corresponding drop-down list, and click **Next**.



- 2 Specify the service to be monitored or a user-defined script to monitor a service.



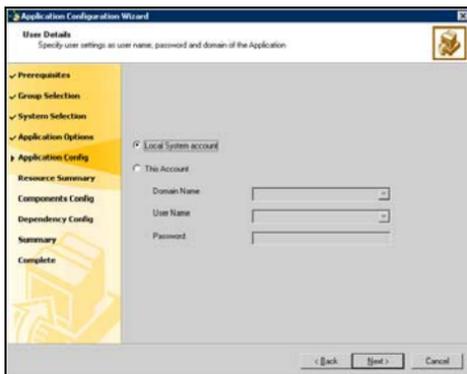
If you want VCS to monitor the service:

- Select the **Service** option and click the icon (...) adjacent to the **Service Name** text box.
- In the Service dialog box, select the service and click **OK**. The selected service name appears in the **Service Name** text box. Alternatively, You may also type in the service name to be monitored.
- Click **Next**.

If you want a script to monitor the service:

- Specify the complete path for the script using the Browse button (...).
- Specify the parameters for the script.
- Specify the time in seconds for the agent to receive a return value from the monitor script.
- Click **Next**.

- 3 Specify the user information in whose context the service will be monitored.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.

If the service selected in [step 2](#) on page 168 is running in the context of a local system account, the **This Account** option is disabled. Similarly, if

the service is running in the context of any other user account, the **Local System account** option is disabled.

- Click **Next**.

ServiceMonitor resource belongs to the category of *persistence* resources. Such resources do not depend on other VCS resources, including shared storage. Hence, the Shared Storage Option panel does not appear if you select the ServiceMonitor option.

- 4 In the Application Resource Summary panel, review the summary of the ServiceMonitor resource. Click **Back** to make changes. Otherwise, click **Next**.
- 5 The Application Options panel appears. Select one of the following options:
 - To configure another ServiceMonitor resource, repeat [step 1](#) through [step 4](#).
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 160 for instructions.
 - To configure a Process resource, see [“Configuring processes”](#) on page 164 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to [“Configuring VCS components”](#) on page 171 for instructions.

If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 174.

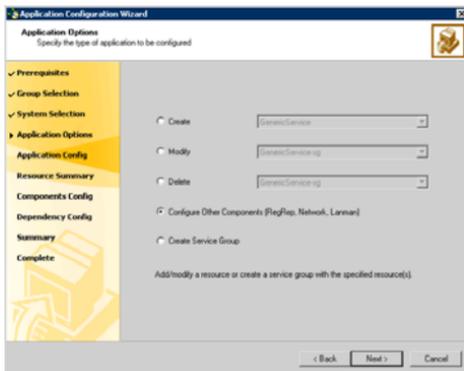
Configuring VCS components

Applications configured using GenericService or Process resources may require network components, or Registry Replication resources. You can configure these VCS components *only* for service groups created using the wizard.

Note: Configure these components only after configuring all application resources. The wizard creates a service group after these components are configured. To add more application resources, you must rerun the wizard in the Modify mode.

To configure VCS components

- 1 In the Application Options panel, click **Configure Other Components**.



- 2 Select the VCS component to be configured for your applications.
The available options are:

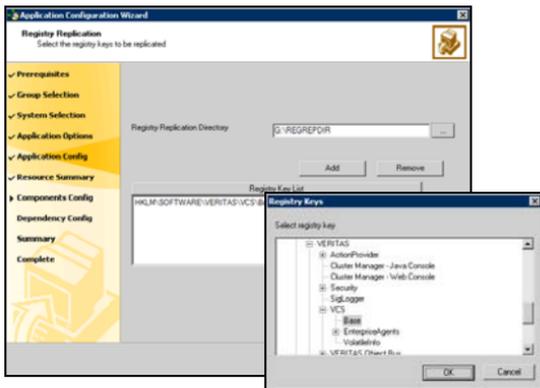
- **Registry Replication Component:** Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to [step 3](#) on page 172.
- **Network Component:** Select this option to configure network components for your application. If you wish to configure a virtual computer name, check **Lanman component** also. To configure a network resource, proceed to [step 4](#) on page 173.

The wizard does not enable the **Lanman Component** check box unless the **Network Component** check box is checked.

To configure Registry Replication

The RegistryReplication panel appears only if you chose to configure the Registry Replication Component in the Application Component panel.

- 3 Specify the registry keys to be replicated.



- Specify the directory on the shared disk in which the registry changes are logged.
- Click **Add**.
- In the Registry Keys dialog box, select the registry key to be replicated.
- Click **OK**.
- The selected registry key is added to Registry KeyList box. Click **Next**.

If you chose Network Component from the Application Component panel, proceed to the next step. Otherwise, proceed to [step 5](#) on page 173.

To configure network components

The Virtual Computer Configuration panel appears only if you chose to configure the Network Component in the Application Component panel.

- 4 Specify information related to your network.
 - Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 16 characters.

Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component panel.
 - Enter a unique virtual IP address for the virtual server.
 - Enter the subnet to which the virtual server belongs.
 - Click **Advanced** to specify additional details for the Lanman resource.
 - Check **AD Update required** to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
 - In the Organizational Unit field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
 - For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

Note that the wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private.
 - Click **Next**.
- 5 The Application Options panel is displayed. Select one of the following options:

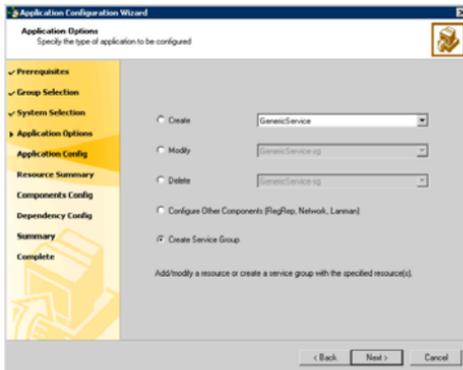
- To configure additional VCS components, repeat [step 1](#) on page 171 through [step 4](#) on page 173.
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 160 for instructions.
 - To configure a Process resource, see [“Configuring processes”](#) on page 164 for instructions.
 - To configure a Service Monitor resource, see [“Configuring a ServiceMonitor resource”](#) on page 168 for instructions.
- If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 174.

Configuring Application Dependencies

The Application Configuration Wizard enables you to create service group for the application resources and other VCS components configured using the wizard. This section describes how to create the service group using the wizard.

To create a service group

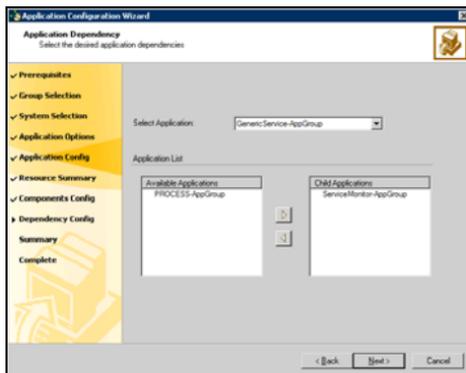
- 1 In the Application Options panel, click **Configure application dependency and create service group**.



The option is enabled only if:

- resources and VCS components are already configured using the wizard.
- you clicked **Modify Service Groups** in the Wizard Options panel.

2 Specify the dependency between the applications.

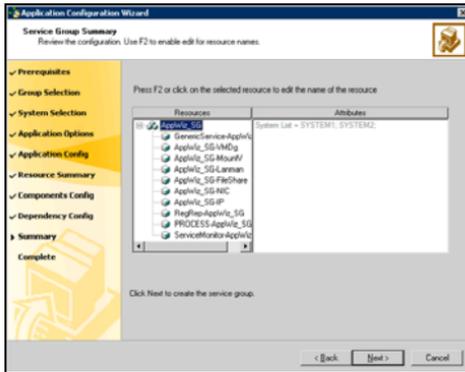


You must have at least two resources configured for Application Dependency panel to appear. Of the two resources, one should either be a GenericService or a Process resource.

- From the Select Application list, select the application that would depend on other applications. The selected application becomes the parent application.
- From the Available Applications list, select the application on which the parent application would depend and click the right-arrow icon to move the application to the Child Applications list.
- To remove an application from the Child Applications list, select the application in the list and click the left arrow.
- Repeat these steps for all such applications for which you want to create a dependency.
- Click **Next**.

The Application Dependency panel enables you to link resources configured using the wizard. If these resources are dependent on other services outside the VCS environment, you should first configure resources for such services and then create the appropriate dependency.

3 Review the service group configuration.



The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

- The wizard assigns unique names to resources. Change names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.
 - Click **Next**.
 - A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.
The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 4 In the completion panel, check **Bring the service group online** if you want to bring the service group online on the local system.
 - 5 Click **Finish** to create the service group and exit the Application Configuration Wizard.

Configuring an Oracle service group

Configuring the Oracle database agent involves creating the Oracle service group, its resources, and defining attribute values for the configured resources. VCS provides several ways to configure the Oracle agent: the agent configuration wizard, the Java and Web consoles, and the command line. This section provides instructions on how to use the agent configuration wizard to configure the agent.

Prerequisites

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
 - You must be a Domain Administrator on the node where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Verify that SFW HA, along with the VCS database agent for Oracle, is installed on all cluster nodes.
 - Verify a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
 - Verify that the Veritas high availability engine (HAD) is running on the system from where you run the wizard.
 - Mount the shared drives containing the data files, control files, redo log files, bdump, cdump, and udump files. Unmount the drives from other systems in the cluster.
 - If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
 - Assign the virtual IP address on the system where you run the wizard. Remove the virtual IP address from other systems

- Mount the database and start the Oracle instance on the node running the wizard.
- Make sure that you have the following information ready; the wizard will prompt you for this information:
 - The databases and listeners to be monitored by VCS.
 - For the instances to be monitored in detail, name and location of the respective SQL files.
 - A valid domain name, user name, and password with which the database service was configured for the database.

Creating an Oracle service group

This section describes how to create an Oracle service group using the Oracle Agent Configuration Wizard.

To create an Oracle service group

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Oracle Agent Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Agent Configuration Wizard**.
- 2 In the Welcome panel click **Next**.
- 3 In the Wizard Options panel, select **Create service group** and click **Next**.
- 4 In the Service Group Configuration panel, complete the following and click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

Service Group Name

Type a name for the Oracle service group.

Available Cluster Systems Select the systems on which to configure the service group and click the right arrow to move the systems to the Systems in Priority Order box.

The Systems in Priority Order box represents the service group's system list. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

To remove a system from the service group's system list, select a system in the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, select the system from the Systems in Priority Order box, and click the up and down arrows.

- 5 In the Oracle Configuration panel, select the SIDs and the associated listeners to be added to the service group, and click **Next**.
The SID is a system identifier that uniquely identifies the Oracle database instance, while the listener is the name of the corresponding listener service.

- 6 On the Detail Monitoring panel, configure detail monitoring for the Oracle database if required, and click **Next**.

Detail Monitor Check the Detail Monitor option corresponding to each database that you want to configure detail monitoring for.

SQL Path Type the path of the SQL file that will query the database to validate the status. Click the icon next to the field to browse for the SQL file.

A sample SQL file, check.sql, is located at
%VCS_HOME%\bin\Oracle\.

- 7 In the Domain and User selection panel, type a valid domain name, user name, and password with which the database service was configured for the database and click **Next**.
- 8 On the Network Configuration panel, specify the network related information and click **Next**.
The wizard discovers and displays the virtual IP address for the Oracle Server.

Subnet Mask Type the subnet mask to which the virtual IP belongs.

Adapter Display Name For each system in the cluster, select the public network adapter. Select the Adapter Name field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- 9 Review the configuration on the Summary panel and click **Next**.

Resources Lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box. The wizard assigns unique names to the resources. To edit a resource name, select the resource name and click on it, or press the F2 key. After the edit, press the Esc key to cancel the changes, or press the Enter key to confirm the changes.

- 10 In the confirmation dialog box, click **Yes**. Click **No** if you wish to review your settings.
The wizard starts running commands to create the Oracle service group.
- 11 On the Completing the Oracle Configuration panel, check **Bring the service group online** to bring the service group online on the local system, and click **Finish**. The Oracle service group is created in your cluster.

Configuring dependent services

If the database service has other dependent services, make sure the dependent services are running on the node where the database service is online. Note that the online agent operation brings only the database service online and not the dependent services.

For example, on Oracle 10g, the DBConsole service corresponding to an Oracle database has a dependency on the database service. That is, for the Enterprise Manager to manage the databases, you must make sure the DBConsole service is running on the node where the database service is online.

To configure a dependent service

- 1 For the dependent service, add a GenericService resource manually.
- 2 Make the GenericService resource dependent on the corresponding Oracle resource.
- 3 Set the Critical attribute to False if the Oracle service group must not fail over when the GenericService resource faults.

Refer to the *Veritas Cluster Server Administrator's Guide* for configuration instructions.

Verifying the cluster configuration

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in [step 1](#).
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.

The service group you selected is taken offline and brought online on the node that you selected.

Possible tasks after completing the configuration

After completing the configuration, you may want to make some changes to the cluster configuration or modify the application service groups. Depending on your specific requirements perform one of the following operations:

- Configuring the Cluster Management Console to centrally administer multiple clusters.
See “[Configuring the Cluster Management Console connection](#)” on page 183.
- Modifying the existing cluster configuration to add additional nodes, removing existing nodes, or configuring the Web Console.
See “[Modifying the existing cluster configuration](#)” on page 188.
- Modifying the replication service group configuration.
See “[Modifying the application service groups](#)” on page 193.

Configuring the Cluster Management Console connection

The Veritas Cluster Management Console (CMC) is a centralized management solution for high-availability application environments based on Veritas Cluster Server. CMC can be configured to locally manage a single cluster or to centrally manage multiple clusters.

CMC comprises of the following components:

- **Management Server**
The management server accepts and processes the operational commands and the configuration inputs that users enter through CMC. The management server communicates with the VCS High Availability engine (HAD). Install the CMC Management Server only if you plan to centrally manage multiple clusters. You must install the management server on a standalone system that is outside any cluster but available on the local network.
- **Cluster Connector**
The cluster connector is an agent that enables the management server to communicate with clusters through intervening firewalls. You must install the cluster connector on each cluster that is separated from the management server by a firewall. If there are no firewalls between the management server and the clusters, you can configure the clusters to use direct connection instead.
In each cluster, the cluster connector runs on one node at a time, but is installed on all nodes and is configured for failover.

This section describes how to install the cluster connector on VCS clusters. For more information on CMC and its components, see the *Veritas Cluster Management Console Implementation Guide*.

Prerequisites for installing the cluster connector

- You must stop all VCS Web consoles, VCS Java consoles, and agent wizards that are running on any cluster nodes before you install the cluster connector
- When you install the cluster connector, Symantec Product Authentication Service must be available on the system from which you run the installer. If you install from a standalone system, you must manually install the authentication service on that system before you install the cluster connector. If you install from a cluster node that is also a member of the target cluster, the installer provides the authentication service automatically.

Possible tasks after completing the configuration

- When installing the cluster connector on 64-bit Windows platforms from a 32-bit system, the default installation directory is C:\Program Files. Symantec recommends that you change the 64-bit installation directory to C:\Program Files (x86).
- Ensure that your network and DNS configuration provide proper name resolution. Otherwise, the cluster connector cannot resolve the management server host name when attempting to connect to the management server.
- The cluster connector requires the management server network address. For example, mgmtserver1.symantecexample.com.
- A CMC service account password. You must have set this account password while installing the management server.
- The root hash of the management server. Use the `vssat showbrokerhash` command and copy the root hash of the management server. Note that you must run this command from the C:\Program Files\Veritas\Security\Authentication\bin directory on the management server.
- After you install and configure the cluster connector, configure the CMC group on all the nodes in the cluster, and the state of the CMC group should ONLINE on one of the cluster nodes.

Installing the cluster connector on Windows clusters

Perform this procedure to use the cluster connector for management server communications with a supported Windows cluster.

To install the cluster connector on a Windows cluster

- 1 Start the Setup program to install the Cluster Connector for Windows.
- 2 In the Symantec Product Installer window, select **VCSMC Cluster Connector for Windows** to install the cluster connector.
- 3 In the Welcome dialog box, make sure all the prerequisites for installing the VCS MC Cluster Connector 5.1 for Windows are satisfied. Click **Next**.
- 4 In the VCS MC Cluster Connector 5.1 for Windows dialog box:
 - Select the domain name and the nodes on which the cluster connector will be installed. Click **Add**.
 - To change the install path, click **Change**.
 - Click **Next**.
- 5 The installer validates the selected nodes in the Validation dialog box. The installation proceeds only if all the nodes are accepted. Click **Next**.

- 6 The installer displays a summary of install options prior to the actual installation. Click **Next**.
- 7 The installation starts on all nodes simultaneously.
- 8 The installer displays the installation report after the installation is completed on all the nodes. Click **Next**.

Click **View Log Files** to see the log files of the installation process. You can check the ClusterConnector-0.log at the following path: C:\Program Files\Symantec\VRTScmccc\log

Check the ClusterConnectorConfig-0.log in the same directory for the cluster connector configuration process.

Configuring the cluster connector

Perform the following steps to configure the cluster connector.

To configure the cluster connector

- 1 Install the management server and configure it. Refer to the *Veritas Cluster Management Console Implementation Guide*.
- 2 Install the cluster connector on a VCS cluster.
- 3 Run the cluster connector configuration utility, found in X:\Program Files\Symantec\VRTScmccc\bin\cc_configure.bat (where X is the driver letter on which the cluster connector is installed).
- 4 Enter the network IP address of the management server or the hostname.
- 5 Enter the certificate to add to the trusted keystore or enter 'q' to quit.
- 6 Enter an administrator user name: **root**
- 7 Enter the domain name. For example **vcs01.symantecexample.com**
- 8 Enter the domain type:
1: Windows
2: nis
3: nisplus
4: unixpwd
5: ldap
6: localhost
0: Quit
Enter the domain type [1]: 4
- 9 Enter the password.
- 10 Enter a unique identifier for the cluster:

Possible tasks after completing the configuration

```
Enter a unique identifier for the
cluster: [43896e6c-0220-4832-9556-97082515c77b]/accept
default:
```

This indicates the configuration is successful.

- 11 To verify that the CMC group and its resources are fully-functional i.e. they are online, can fail over, etc., check for the existence of the cluster on the management server.

Configuring the cluster connector using the management server console

This task enables you to configure an upgraded version of the cluster connector. Before you perform this task, you must first install an upgraded version of the cluster connector on the target clusters. This task configures only versions of the cluster connector that have already been installed on the target clusters.

To upgrade the cluster connector on discovered clusters

- 1 On the main tab bar, click **Administration**.
- 2 On the details tab bar, click **Configured Clusters**.
- 3 In the Configured Clusters table, do one of the following:
 - To select one or more clusters, check the check box next to each required cluster.
 - To select all clusters, check the check box at the top of the table.
- 4 On the Configuration task menu, select **Upgrade Cluster Connector**.
- 5 In the Upgrade Cluster Connector wizard, read the overview information and then click **Next**.
- 6 This launches the **Upgrade Cluster Connector** wizard to configure known (secure or non-secure clusters). Click **Next**.
- 7 In the Access Credentials for Target Clusters panel, specify the following options:
 - The type of security access that the cluster uses. The options are:
 - Classic VCS
This option enables only VCS users that are configured locally on this cluster to log in to the cluster.
 - VxAT
Otherwise known as Symantec Product Authentication Service, VxAT is the Symantec cross-product user authentication service. If you select VxAT, you must also specify the IP address of the Symantec Product authentication broker that you want to use.

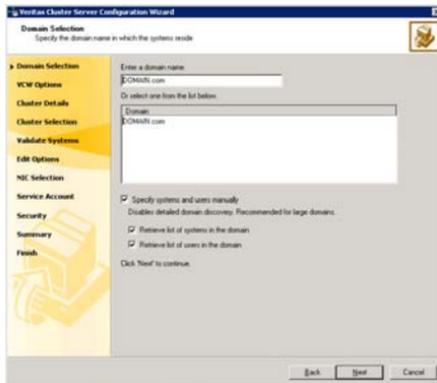
- The cluster administrator user name, password, domain, and domain type required to establish a connection to the cluster. You must be a cluster-level administrator on each cluster that you want to add or discover. The **Domain** field requires a fully qualified domain name.
- 8 To configure clusters in the secure mode in the Discover Clusters dialog box:
 - Select **VxAT**.
 - Enter the access credentials (user name and password) of the target clusters.
 - Click **Next**.
 - 9 To configure clusters in the non-secure mode in the Discover Clusters dialog box:
 - Select **Classic VCS**.
 - Enter the access credentials (user name and password) of the target clusters.
 - Click **Next**.

If you have specified both VxAT security clusters and Classic VCS security clusters, this panel runs separately for each. The wizard enables you to select either the cluster's authentication broker or one of the predefined authentication brokers.
 - 10 In the Summary of Target Clusters panel, read the overview of your selections and then click **Finish**.

Modifying the existing cluster configuration

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



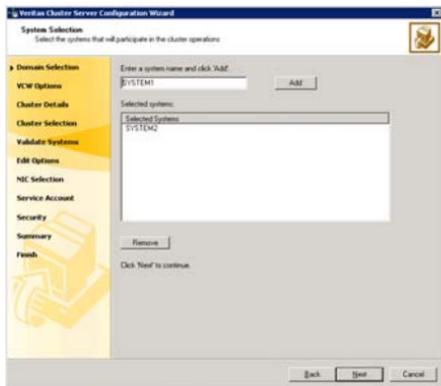
Do one of the following:

- To discover information about all the systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 191.

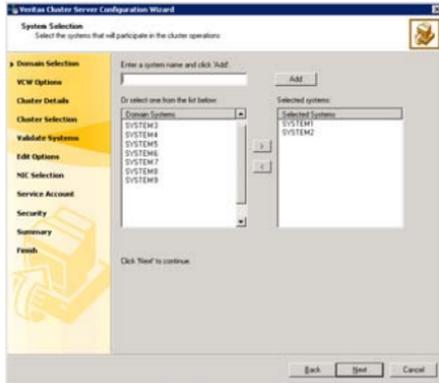
- To specify systems and user names manually (recommended for large domains):

- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 190. Otherwise proceed to the next step.
- 5 On the System Selection panel, complete the following and click **Next**.



- Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 8](#) on page 191.

- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.



Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

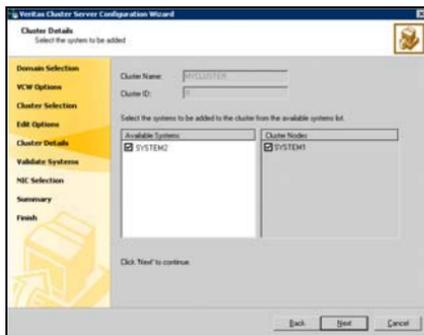
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**. In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.



The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on

how it is configured in the cluster. If LLT is configured over ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have to use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard will configure the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select **Configure LLT over UDP**.
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 14 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
- 17 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

Modifying the application service groups

You may want to modify existing application service groups. Use one of the following options depending on your specific application environment:

- [Modifying the FileShare service group](#)
- [Modifying the PrintShare service group](#)
- [Modifying the IIS service group](#)
- [Modifying the MSVirtual Machine service group](#)
- [Modifying any other application service group](#)
- [Modifying an Oracle service group](#)

Modifying the FileShare service group

To modify a File Share service group

- 1 Start the File Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > File Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, select **Modify service group**, select the service group to be modified, and click **Next**.
- 4 Follow the wizard instructions and make required modifications to the service group configuration.

Modifying the PrintShare service group

The Print Share Configuration Wizard enables you to modify a Print Share service group.

Following are some points to note before modifying the service group:

- If the Print Share service group is online, you must run the wizard from a system on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot modify resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.
- If the service group contains resources that were not part of the default service group configuration, then modifying the service group may change those resources. You may then have to manually restore the settings of those resources later.

To modify the Print Share service group

- 1 Start the Print Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.

From the Solutions Configurations Center expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.

- 2 Read the information on the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 4 Follow the wizard instructions and make required modifications to the service group configuration.
If you are modifying the service group to remove a PrintShare resource, make sure you offline the resource before deleting it.

Modifying the IIS service group

The IIS configuration wizard enables you to modify an IIS service group.

Note the following before you modify an IIS service group:

- If the IIS service group is online, you must run the wizard from a system on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot modify resources that are online.
- To change the online resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.
- If the service group contains resources that were not part of the default service group configuration, then modifying the service group may change those resources. You may then have to manually restore the settings of those resources later.

To modify the IIS service group

- 1 Start the IIS Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- 2 From the **Solutions Configurations Center**, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > IIS Configuration Wizard**.
- 3 Review the information in the Welcome panel and click **Next**.

Possible tasks after completing the configuration

- 4 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 5 Follow the wizard instructions and make required modifications to the service group configuration.

Modifying the MSVirtual Machine service group

To modify the MSVirtual Machine service group

- 1 Start the MSVirtual Machine Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > MSVirtual Machine Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, choose the **Modify service group** option, select the service group to be modified, and click **Next**.
- 4 Follow the wizard instructions and make required modifications to the service group configuration.

Modifying any other application service group

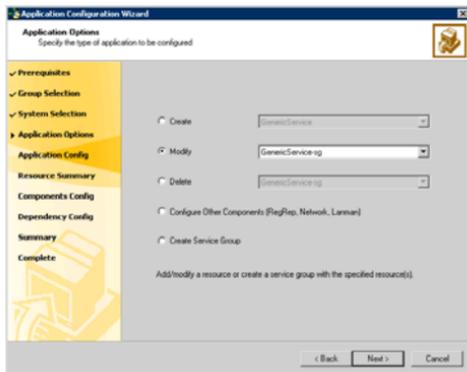
This section describes how to modify a service group using the Application Configuration Wizard. Following are some points to note before modifying the service group:

- If the service group to be modified is online, you must run the wizard from a system on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot modify resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.
- If the service group contains resources that were not part of the default service group configuration, then modifying the service group may change those resources. You may then have to manually restore the settings of those resources later.

Note: Symantec recommends that you do not use the wizard to modify service groups that were not created using the wizard.

To modify a service group

- 1 Start the Application Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the Solutions Configurations Center, expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Application Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**. From the Service Groups list, select the service group containing the resource that you want to modify.
- 4 In the Service Group Configuration panel, click **Next**.
- 5 Click **Modify** and select the resource you want to modify and then click **Next**.



The **Modify** option is enabled only if:

- Service and Process resources are already configured using the wizard.
 - You selected the **Modify Service Groups** option in the **Wizard Options** dialog box.
- 6 Depending on the resource you chose to modify from the Application Options page, you would either get the Generic Service Options, Process

Details, or Service Monitor Options dialog box. Make required changes in the appropriate dialog box and click **Next**.

- 7 In the User Details dialog box, specify the user information and click **Next**.
- 8 In the Application Resource Summary dialog box, review the summary of the resource. Click **Back** to make changes. Otherwise, click **Next**.
- 9 The Application Options dialog box appears. Repeat [step 5](#) through [step 8](#) for each resource that you want to modify.
- 10 After modifying the required resources, you can use the wizard to:
 - Add additional resources to the service group.
 - Delete resources from the service group.
 - Add VCS components to the service group.

Modifying an Oracle service group

The following steps describe how to modify the configuration of the service groups using the configuration wizard.

Prerequisites

- If the Oracle service group is online, you must run the wizard from a node on which the service group is online. You can then add and remove resources to the configuration using the wizard; you cannot modify resources that are online.
You can however enable and disable the SIDs and listeners to be monitored in detail and change their detail monitoring options when the Oracle service group is online.
- To change the online resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.

Instructions

To modify an Oracle service group

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA)**

Configuration > Configure the Service Group > Oracle Agent Configuration Wizard.

- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group to modify, and click **Next**.
- 4 Follow the wizard instructions and make desired modifications to the service group configuration.

Deleting an Oracle service group

The following steps describe how to delete an Oracle service group using the configuration wizard.

To delete an Oracle service group

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Oracle Agent Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Delete service group**, select the service group to delete, and click **Next**.
- 4 In the Service Group Summary panel, click **Next**.
- 5 On the message that informs you that the wizard will run commands to delete the service group, click **Yes** to delete the service group and then click **Finish**.

Configuring detail monitoring

Use the detail monitoring capability of VCS database agent for Oracle to monitor the status of a database. Before setting up detail monitoring, you must have the agent running at the basic level of monitoring, that is, the DetailMonitor attribute must be set to 0.

The Oracle agent uses a script to monitor the status of the database. A sample SQL script, located at %VCS_HOME%\bin\Oracle\check.sql, is provided with the agent for the purpose. If the script is successfully executed during monitoring, the agent considers the database instance available. If the execution fails, the database instance is considered not available and the service group

faults and fails over to the failover nodes. You can customize the script to meet your configuration requirements.

Note: You must use a separate script for each Oracle service group that exists in the cluster. The script must exist on all the nodes in the service group.

Enabling detail monitoring

The following steps describe how to enable detail monitoring using the configuration wizard.

To enable detail monitoring

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Oracle Agent Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group configured for the instance to be monitored in detail, and click **Next**.
- 4 In the Oracle Configuration panel, select the SIDs to be configured along with their respective listeners and click **Next**.
- 5 In the Detail Monitoring dialog box, specify information to enable detail monitoring, and click **Next**.

Detail Monitor Check the Detail Monitor option corresponding to each database that you want to configure detail monitoring for.

SQL Path Type the path of the SQL file that will query the database to validate the status. Click the icon next to the field to browse for the SQL file.
A sample SQL file, check.sql, is located at
%VCS_HOME%\bin\Oracle\.

- 6 In the Domain and User selection panel, type a valid domain name, user name, password, and then click **Next**.
- 7 Follow the wizard instructions and accept the default values in the subsequent dialog boxes.

Disabling detail monitoring

The following steps describe how to disable detail monitoring using the configuration wizard.

To disable detail monitoring

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **High Availability (HA) Configuration > Configure the Service Group > Oracle Agent Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Wizard Options panel, click **Modify service group**, select the service group configured for the instance for which detail monitoring is being disabled, and click **Next**.
- 4 In the Oracle Configuration panel, click **Next**.
- 5 In the Detail Monitoring Configuration panel, uncheck the check box corresponding to the Oracle Server instance for which detail monitoring is being disabled and click **Next**.
- 6 Follow the wizard instructions and accept the default values in the subsequent dialog boxes.

Campus Clustering

The section includes the following chapters:

- [Introduction to campus clustering](#)
- [Deploying SFW HA for campus cluster](#)

Introduction to campus clustering

This chapter contains the following topics:

- [“About Campus Clusters”](#) on page 206
- [“Sample campus cluster configuration”](#) on page 206
- [“Differences between campus clusters and local clusters”](#) on page 208

About Campus Clusters

A campus cluster is a single cluster that stretches over two sites using fiber channel connectivity, with SAN connections for data mirroring and network connections for cluster communication. Although two sites are the most common, more than two can be used for additional redundancy.

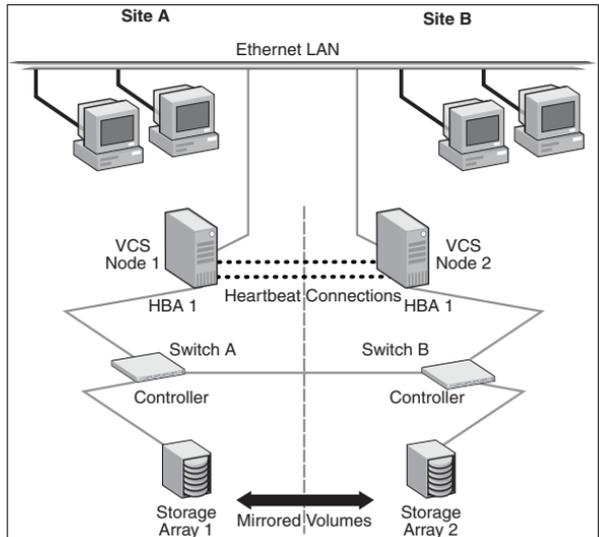
Clusters are usually located across a campus or a city but can range over much wider distances if their infrastructure supports it, using Fibre Channel SANs and long-wave optical technologies.

Storage administrators can protect their clusters by using campus clusters to protect from natural disasters, such as floods and hurricanes, and from unpredictable power blackouts. Campus clusters provide disaster protection when an entire site goes down by locating the clustered servers in different buildings or areas. This solution provides a level of high availability that is above mirroring or clustering at a single site and is an alternative to using replication software.

Sample campus cluster configuration

The following sample configuration represents a campus cluster with two sites, Site A and Site B.

Figure 9-1 Typical campus clustering configuration



With SFW, a campus cluster can be set up using a Veritas Cluster Server (VCS) configuration. Both configurations involve setting up a single cluster with two nodes that are in separate buildings and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. SFW provides the mirrored storage and the disk groups that make it possible to fail over the storage by deporting the disk groups on one node and importing them on the other.

If a site failure occurs in a two-node campus cluster, the remaining cluster node will not be able to bring the cluster disk groups online because it cannot reserve a majority of disks in the disk groups. To allow for failover to the other site, a procedure forces the import to the other node, allowing a cluster disk group to be brought online on another node when that node has a minority of the cluster disks.

Implementing these force import procedures should be done with care. The primary site may appear to have failed but what really has happened is that both the storage interconnect between sites and the heartbeats have been lost. In that

case, cluster disk groups can still be online on the primary node. If a force import is done so that the data can be accessed on the secondary site, the cluster disks will be online on both sites, risking data corruption.

Differences between campus clusters and local clusters

The procedures for setting up a campus cluster are nearly the same as those for local clusters, except that a campus cluster has the nodes located in separate buildings, so the hardware setup requires SAN interconnects that allows these connections. Also, in a campus cluster, each node has its own storage array rather than having a shared storage array between the two clusters. Both local clusters and campus clusters have SFW dynamic disk groups and volumes, but the volumes on each campus cluster node are mirrors of one another.

Deploying SFW HA for campus cluster

This chapter contains the following topics:

- [“About the Campus Cluster solution”](#) on page 210
- [“Reviewing the requirements”](#) on page 212
- [“Reviewing the configuration”](#) on page 219
- [“Installing and configuring the hardware”](#) on page 224
- [“Configuring the storage hardware and network”](#) on page 225
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 228
- [“Resetting the driver signing options”](#) on page 233
- [“Configuring the cluster”](#) on page 233
- [“Creating disk groups and volumes”](#) on page 249
- [“Installing the application on cluster nodes”](#) on page 257
- [“Configuring service groups”](#) on page 261
- [“Verifying the cluster configuration”](#) on page 320

About the Campus Cluster solution

This chapter presents a VCS campus clustering configuration example with SFW HA.

For a campus clustering configuration with Microsoft clustering, see the following:

- Chapter 19, “Deploying SFW with MSCS in a campus cluster” on page 675.
- Chapter 20, “Deploying SFW with Microsoft failover clustering in a campus cluster” on page 717

The table below outlines the configuration’s high-level objectives and the tasks for each objective.

Table 10-1 Task list: Campus Cluster configuration

Objectives	Tasks
“Reviewing the requirements” on page 212	Verify hardware and software prerequisites.
“Reviewing the configuration” on page 219	<ul style="list-style-type: none">■ Review the configuration requirements.■ Overview of VCS campus cluster, and recovery scenarios
“Installing and configuring the hardware” on page 224	<ul style="list-style-type: none">■ Install the hardware for Site A.■ Install the hardware in the same manner for Site B.■ Make all the necessary hardware connections between the two cluster nodes.
“Configuring the storage hardware and network” on page 225	<ul style="list-style-type: none">■ Install the operating system on both nodes.■ Make necessary networking settings on both nodes.
“Installing Veritas Storage Foundation HA for Windows” on page 228	<ul style="list-style-type: none">■ Set the Windows driver signing options to “Ignore” for remote nodes if you are using Windows Server 2003.■ Install SFW HA on both nodes with the push install. Doing so installs SFW, VCS, and the Veritas Cluster Server Enterprise Agent on both cluster nodes.
“Configuring the cluster” on page 233	Use the VCS Cluster Configuration Wizard (VCW) to set up the cluster.

Table 10-1 Task list: Campus Cluster configuration (Continued)

Objectives	Tasks
“Creating disk groups and volumes” on page 249	<ul style="list-style-type: none">■ Create dynamic cluster disk groups.■ Create dynamic volumes.
“Installing the application on cluster nodes” on page 257	<ul style="list-style-type: none">■ Install the application program files on the local drive of the first node.■ Install files relating to the data and logs on the shared storage.■ Deport the disk groups on the first node and import them on the second node.■ Install the application on the second node.
“Configuring service groups” on page 261	<ul style="list-style-type: none">■ Use an appropriate method to create and configure the VCS service group or groups.■ Bring the service group online.
“Verifying the cluster configuration” on page 320	<ul style="list-style-type: none">■ Switch the service group to the second node.■ Switch it back to the first node.

Reviewing the requirements

Review these product installation requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

Note: Before you install the SFW or SFW HA software, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware compatibility List to confirm supported hardware:
<http://www.symantec.com/business/support/index.jsp>

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space.

[Table 10-2](#) estimates disk space requirements for SFW HA.

Table 10-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 5.1 Service Pack 2 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/358407>
- Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:

<http://entsupport.symantec.com/docs/358406>

- Review the operating systems supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported operating systems

SFW and SFW HA have client and server components that run on specific Windows operating systems.

SFW and SFW HA software for servers

Your system must run one of the operating systems listed in the tables to install the SFW or SFW HA server software:

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 10-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required

Table 10-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

The following table shows supported Windows Server 2008 operating systems.

Note: SFW HA supports Windows Server 2008 R2 without Hyper-V x64: Standard Edition, Enterprise Edition, Datacenter Edition.

Table 10-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server Note: SFW HA does not support Hyper-V.	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition Note: SFW HA does not support Hyper-V.	
Windows Server 2008 R2 for IA systems	IA64		

Table 10-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Storage Server 2008			

SFW and SFW HA software for clients

The SFW or SFW HA client software can be installed on any one of the operating system versions, editions, and architectures that the Server Components are supported on, except Server Core.

The following table shows the additional client operating systems that are supported.

Table 10-5 Supported client operating systems

Windows Client	Platform	Edition	Service Pack
Windows XP	x86, x64		SP3 required
Windows Vista	x86, x64,	Ultimate Edition, Business Edition, Premium Edition	SP1 or SP2 required
Windows 7	x86, x64	Ultimate Edition, Business Edition, Premium Edition	

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs. See "[Best practices for SFW HA](#)" on page 218.
- NIC teaming is not supported for the VCS private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each application virtual server.
 - A minimum of one static IP address for each physical node in the cluster.

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements for SFW HA

The following permissions are required:

- You must be a domain user.
- You must be a member of the local Administrators group on all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements for SFW HA

Please review the following additional requirements:

- Installation media for all products and third-party applications.
- Licenses for all products and third-party applications.

- You must install the operating system in the same path on all systems. For example, if you install Windows Server on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- For a Replicated Data Cluster, install only in a single domain.

Best practices for SFW HA

Symantec recommends that you perform the following tasks:

- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the vxclus UseSystemBus ON command. This is applicable for a Replicated Data Cluster configuration.

Campus cluster requirements

- Interconnects between the clusters are required for the storage and the network.
- The configuration requires a storage array for each site, with an equal number of disks at each site for the mirrored volumes.

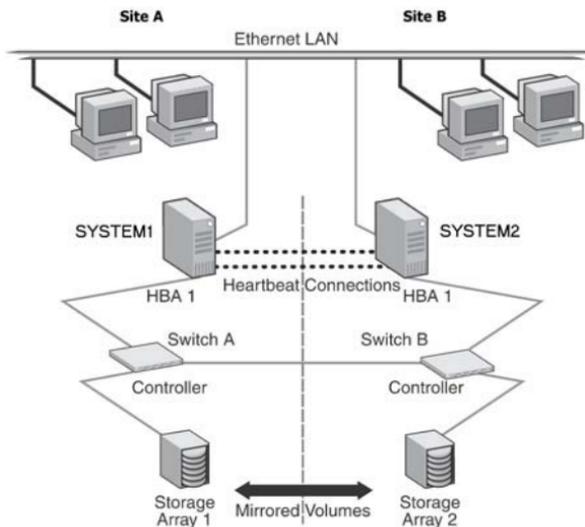
Note: Plan for an equal number of disks on the two sites, because each disk group must contain the same number of disks on each site.

Reviewing the configuration

This configuration example describes the most common configuration, a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with VCS, see [“Overview of campus clustering with VCS”](#) on page 221.

Figure 10-1 VCS 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 with an equal number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

Plan for an equal number and size of disks on the two sites, because each disk group must contain the same number of disks on each site for the mirrored volumes.

The configuration does not include DMP. For instructions on how to add DMP to a clustering configuration, see the DMP chapter, [“Adding DMP to a clustering configuration”](#) on page 67.

Overview of campus clustering with VCS

This overview focuses on the recovery with a VCS campus cluster. Automated recovery is handled differently in a VCS campus cluster than with a VCS local cluster.

The table below lists failure situations and the outcomes that occur with the two different settings for the ForceImport attribute of the VMDg resource. This attribute can be set to 1 (automatically forcing the import of the disk groups to the another node) or 0 (not forcing the import). Information on how to set the ForceImport attribute are given in [“Setting the ForceImport attribute”](#) on page 224.

Table 10-6 Failure Situations

Failure Situation	Forceimport set to 0 (import not forced)	Forceimport set to 1 (automatic force import)
1) Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Application automatically moves to other site.	Service Group failover is automatic on the standby or preferred system or node.
2) Server failure May mean that a power cord was unplugged, a system hang occurred, or another failure caused the system to stop responding.	Application automatically moves to other site. 100% of the disks are still available.	Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Remaining disks in mirror are still accessible from other site.	The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.
4) Site failure All access to the server and storage is lost.	Manual intervention required to move application. Can't import with only 50% of the disks available.	Application automatically moves to the other site.

Table 10-6 Failure Situations (Continued)

Failure Situation	Forceimport set to 0 (import not forced)	Forceimport set to 1 (automatic force import)
<p>5) Split-brain situation (loss of both heartbeats) If the public network link is used as a low-priority heartbeat, it is assumed that link is also lost.</p>	No interruption of service. Can't import disks because original site still has the SCSI reservation.	No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.
<p>6) Storage interconnect lost Fibre interconnect severed.</p>	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	No interruption of service. Service Group resources remain online, but 50% of the mirror disk becomes detached.
<p>7) Split-brain situation and storage interconnect lost If a single pipe is used between buildings for the Ethernet and storage, this situation can occur.</p>	No interruption of service. Can't import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.	Automatically imports disks on secondary site. Now disks are online in both locations—data can be kept from only one.

Reinstating faulted hardware

Once a failure occurs and an application is migrated to another node or site, it is important to know what will happen when the original hardware is reinstated.

For failure scenarios 3 through 7 above, the table below lists the behavior when various hardware components affecting the configuration (array or disks, site hardware, networking cards or cabling, storage interconnect, etc.) are reinstated after failure. Situations 1 and 2 have no effect when reinstated. Keep in mind

that the cluster has already responded to the initial failure as indicated in the table above.

Table 10-7 Behavior exhibited when hardware is reinstated

Failure Situation, before Reinstating the Configuration	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Resync the mirror from the remote site.	Same behavior.
4) Site failure All access to the server and storage is lost.	Inter-node heartbeat communication is restored and the original cluster node becomes aware that the application is online at the remote site. Resync the mirror from the remote site.	Same behavior.
5) Split-brain situation (loss of both heartbeats)	No interruption of service.	Same behavior.
6) Storage interconnect lost Fibre interconnect severed.	No interruption of service. Resync the mirror from the original site.	Same behavior.
7) Split-brain situation and storage interconnect lost	No interruption of service. Resync the mirror from the original site.	VCS alerts administrator that volumes are online at both sites. Resync the mirror from the copy with the latest data.

While the outcomes of using both settings of the ForceImport attribute for most scenarios are the same, the ForceImport option provides automatic failover in the event of site failure. This advantage comes at the cost of potential data loss if all storage and network communication paths between the sites are severed. Choose an option that is suitable given your cluster infrastructure, uptime requirements, and administrative capabilities.

Setting the ForceImport attribute

After the VCS campus cluster is configured, set the ForceImport attribute. The command for implementing the force import setting in VCS is:

```
hares -modify <vmdg_resource_name> ForceImport 1|0
```

ForceImport is a flag that defines whether the agent forcibly imports the disk group when exactly half the disks are available. The value 1 indicates the agent imports the configured disk group when half the disks are available. The value 0 indicates it does not. Default is 0. This means that the disk group will be imported only when SFW acquires control over the majority of the disks.

Caution: Set this attribute to 1 only after verifying the integrity of your data. If due caution is not exercised before setting this attribute to 1, you risk potential data loss.

Example

```
hares -modify vmdg_Dg1 ForceImport 1
```

Import is forced on vmdg_Dg1.

Installing and configuring the hardware

This topic gives the general steps for the hardware installation. For complete details on installing the hardware, refer to the hardware documentation.

To set up the hardware

- 1 Install the hardware for Site A, using the manufacturers' instructions.
 - Install three network interface cards.
 - One for the public network.
 - Two are recommended for the private network.
Use independent hubs or switches for each VCS communication network (GAB and LLT). GAB supports hub-based or switch network paths, or two-system clusters with direct network links.

Note: To prevent lost heartbeats on the private networks and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Symantec also recommends removing TCP/IP from private NICs to lower system overhead. Contact the NIC manufacturer for details on this process.

- Install the host adapter.

- Install the switch and the storage array.
 - Verify that the system can access the storage devices.
- 2 Install the hardware in the same manner for Site B.
 - 3 Make the necessary hardware connections to connect the two clusters together.
 - Connect corresponding cables between the three networking cards on the two sites.
 - Connect the two switches at the two sites through the storage interconnect.
 - 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.

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.

Note: Follow the appropriate procedure for verifying the DNS settings, as the procedures for Windows Server 2003 and Windows Server 2008 are slightly different.

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 VCS 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.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.

- 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.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 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 and Dial-up 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 VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in 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**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 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 VCS 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.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 10-8 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see ["Installing Symantec Trusted certificate for unsigned drivers"](#) on page 229.

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.

Installing Symantec Trusted certificate for unsigned drivers

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 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.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

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 HA, 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, select the appropriate SFW product options, and click **Next**.

Client Components	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used while configuring high availability. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration
Global Cluster Option	Required for a disaster recovery configuration only.
Volume Replicator (VVR)	If you plan to use Veritas Volume Replicator (VVR) for replication, select this option to install VVR.
High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.

- 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.

If you are using multiple paths and selected 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 path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation for Windows (Server Components) with the Microsoft Cluster Server (MSCS) option, 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. For more information, see the *Veritas Storage Foundation Administrator's Guide*.

Review the messages, and then click **OK**.

- 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.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

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 the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management

Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

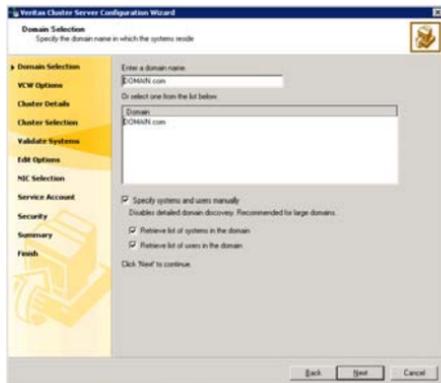
- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

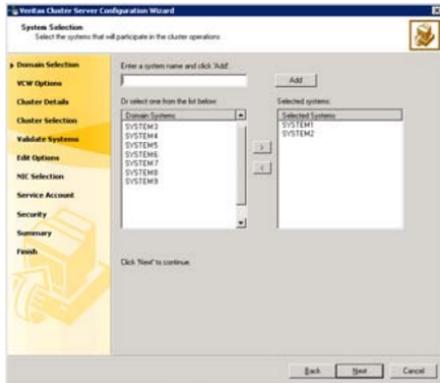
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 236.
 - To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 236. Otherwise, proceed to the next step.
- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster. Proceed to [step 8](#) on page 236.

- 6 On the System Selection panel, specify the systems for the cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

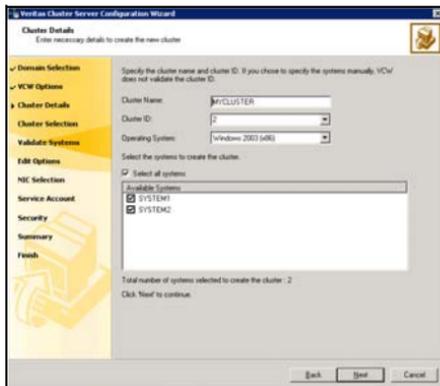
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

Operating System From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (NICs) on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

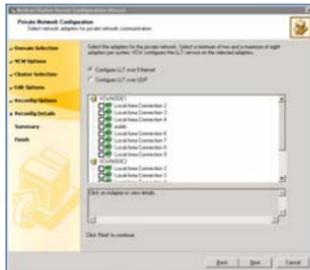
- 10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 240.

- 11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



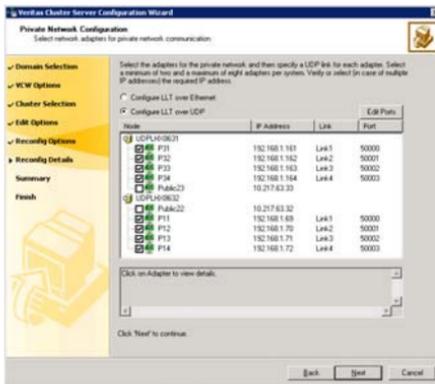
- Select the check boxes next to the two NICs to be assigned to the private network.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links.

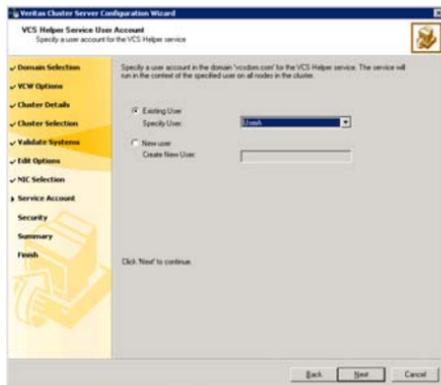
Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

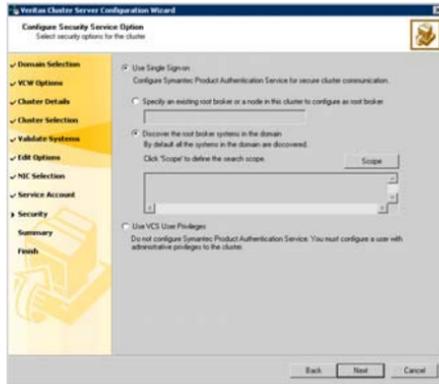
This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 235, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode. For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 10-9 contains some more examples of search criteria.

Table 10-9 Search criteria examples

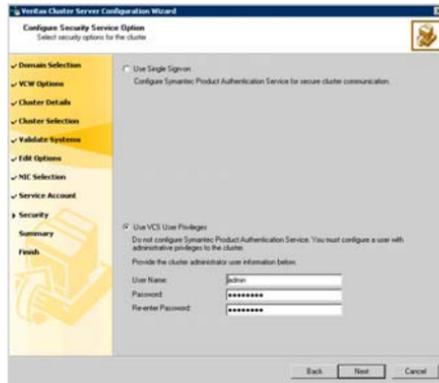
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS Encrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

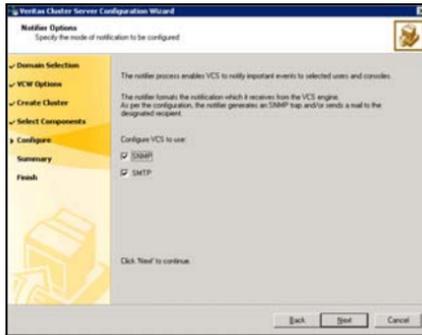
- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 245.

Configuring notification

This section describes steps to configure notification.

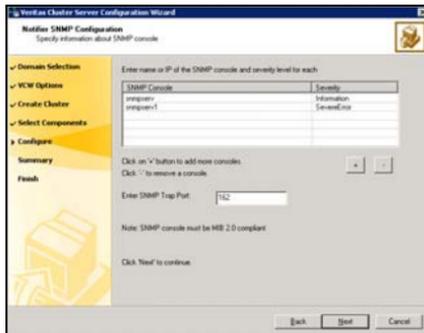
To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

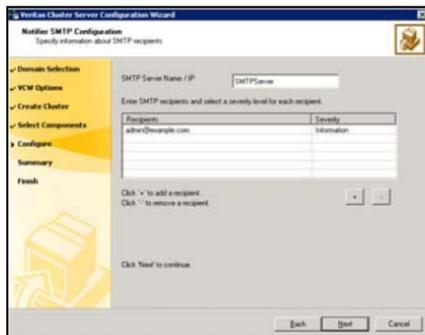


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

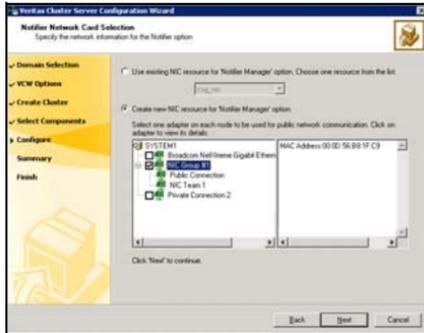


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

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 two storage arrays. Create one or more dynamic cluster disk groups on the storage.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Before you create a disk group, consider the following items:

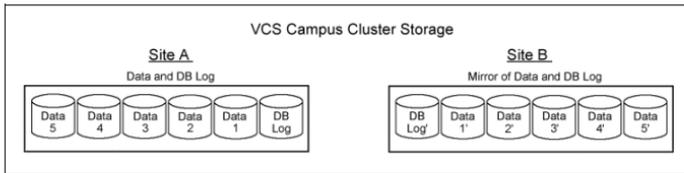
- The type of volume configurations that are required
- The number of LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs that depend on the traffic load
- The disk groups and number of disks on each site
- Types of volumes required and location of the plex of each volume in the storage array

Note: For campus clusters, each disk group must contain an equal number of disks on each site.

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.

The illustration that follows shows a VCS campus cluster configuration of disks. This example has one 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 does not need to be limited to two disks, but can have four disks for greater resiliency. All the data on one site could be in one large mirrored volume with multiple disks, but this also requires the same number of disks on both sites for the mirroring. It is recommended that the log volumes be on separate disks from the data.

Figure 10-2 VCS campus cluster disks example



Configuring the disks and volumes

Ensure that each disk group contains an equal number of disks on each site. Each volume must 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.

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 251
- ["Creating a volume"](#) on page 253

Considerations when creating new volumes

Consider the following when creating new volumes.

- For campus clusters, when creating 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

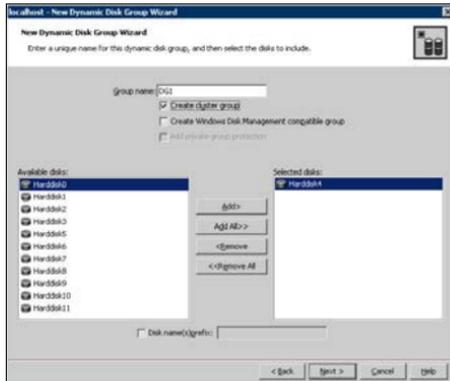
Use the following procedure to create a dynamic disk group.

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 disk group name (for example, DG1).
- 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.

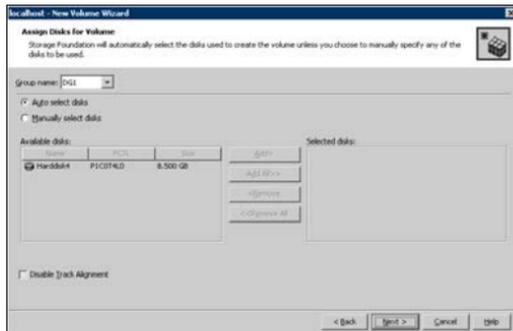
Proceed to create the appropriate volumes on each disk.

Creating a volume

Use the following procedure to create dynamic volumes.

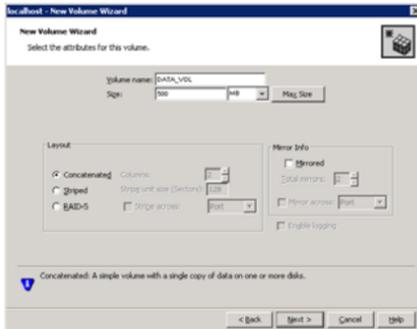
To create dynamic volumes

- 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.
- 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.



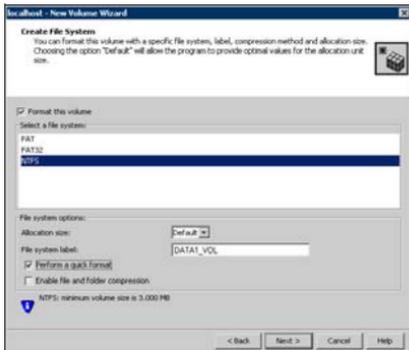
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.
- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that 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 additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Installing the application on cluster nodes

VCS requires that the application program files be installed on the same local drive of all cluster nodes and that the application data and log files or other files related to the application data be installed on the shared storage.

Pointers 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.
- Make sure that the disk groups and volumes are imported and thus mounted on the server before you install the application.
- If you have just created the disk groups and volumes, they will be mounted and accessible. When a disk group is created, it is automatically imported on that node. You can verify that the disk group and volumes are accessible if you can see the disk group and volume icons in VEA for the server.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Remember not to accept the default location for the application data and log files when installing the application. Instead, click to browse to the dynamic volumes that were prepared previously.

Pointers for installing the application on the second node

- To install the application on the second node, deport any disk groups from the first node and import them on the second node. Steps for deporting and importing disk groups are in the section [“Deporting and importing a disk group”](#) on page 260.
- 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. To change a drive letter or mount point, see the instructions in the section [“To add or change a drive letter or mount point”](#) on page 261.
- 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.

Configuring a File Share server role

Points to note when configuring a File Share:

- Make sure that the disk group and volumes for the file server shared directory are configured on shared storage.
- When configuring a new set up, first create the disk groups and volumes on the shared storage and then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.
- The FileShare agent is installed automatically with SFW HA.

Configuring a Print Share server role

Points to note when configuring a Print Share:

- Make sure the printer is connected to the network and is configured with an IP address.
- Install software drivers for the network printer on all systems in the cluster.

To add a print driver

- 1 Open the **Printers** Control Panel.
- 2 Click **File > Server Properties**.
- 3 In the Print Server Properties dialog box, click the **Drivers** tab.
- 4 Click **Add**. This launches the Add Printer Driver wizard.
- 5 Follow the wizard instructions to add the printer driver on the system. You must add the driver on each system that will be part of the service group.

Installing and configuring the IIS application

Points to note when installing IIS:

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on shared storage.
- Import the cluster disk groups and mount the volumes that contain the website data, on the first node.
- For a new IIS installation, while creating new web sites, create the site folder on the shared storage and place the site content in that folder.

- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the website in IIS and then restart the website again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group.

To synchronize the IIS configuration on Windows 2003 systems

- 1 Synchronize the IIS configuration on all nodes that will host the IIS service group. Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system.

For example, the following command copies the IIS metabase to *target_sytem*. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /t  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing and configuring Microsoft Virtual Server

Points to note when installing MS Virtual Server:

- Verify Microsoft Virtual Server is installed and configured identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- Install and configure Virtual Machine Additions *on each virtual machine* if you plan to enable detailed monitoring for the virtual machine resources.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

Installing additional applications

Following are some very generic points for installing any application:

- Make sure that the disk groups and volumes are mounted on the node before installing the application.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes, must be on drive C.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the shared storage.

Deporting and importing a disk group

This section describes the steps for deporting and importing a disk group in order to install the application on the second node.

To deport a disk group on the first node

- 1 If VEA is not already running, start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node on which the dynamic disk group is currently imported.
- 3 Right-click the dynamic disk group to be deported and click **Deport**.

To import the dynamic disk group on the second node

- 1 Start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Enterprise Administrator**). If the Storage Foundation Assistant automatically opens, close it.
- 2 Navigate to **dynamic disk groups** on the node to which you will import the dynamic disk group.
- 3 Right-click the dynamic disk group to be imported and click **Import**.
No drive letter may be associated with an existing dynamic volume when it is imported to a computer for the first time. In such a case, use VEA to add or change drive letters. You need to make sure that drive letters or mount points for the volumes on the second node are the same as were used on the first node.

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**. The Drive Letter and Paths window appears.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Select the new drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

Note: A mount point is also referred to as a “drive path.”

- 6 To change a mount point, you must remove it and recreate it (step 5). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

Configuring service groups

In order for VCS to be able to monitor and fail over an application in a cluster, the application must be included in a VCS service group.

A service group is a collection of resources working together to provide application services to clients. It can also relate to a print or a file share that does not contain a specific application. A service group's resources fail over as a group to another cluster node when there is an application failure or server failure on the active node.

VCS provides multiple methods for creating a service group. If you have Microsoft Exchange Server or SQL Server as the application, VCS provides a wizard for each of these, but you need to purchase the VCS enterprise agents for these programs. There are also separate wizards for file and print servers. In addition, there are several ways to create a service group through VCS Java Console, as well as a generic Application Configuration wizard. If you prefer to use the command line, that method can be used to create a service group as well.

Creating a VCS service group provides the following:

- Defining the cluster resources and their attributes.
- Setting their dependencies; for example, a NIC resource depends on an IP resource.
- Logically grouping the resources together.
- Providing capabilities for monitoring the service group and taking it online or offline.

For an example of installing a service group with the Application Configuration wizard, see the section “[Configuring the service group](#)” on page 124 in [Chapter 8, “Deploying SFW HA for high availability: New installation”](#).

For instructions on how to create a service group for Microsoft Exchange Server or Microsoft SQL Server, see the other Solutions Guides included with this release:

- *Veritas Storage Foundation and High Availability Solutions Solutions Guide for Microsoft Exchange*
- *Veritas Storage Foundation and High Availability Solutions Solutions Guide for Microsoft SQL*

The Solutions Configuration Center provides wizards to configure the service groups for the additional SFW HA applications or server roles. It also supports the Application Configuration Wizard which can be used to configure any other application for which application specific wizards have not been provided. Depending on the application that you have installed, complete the appropriate procedure to configure the service group:

- [Configuring the File Share service group](#)
- [Configuring the Print Share service group](#)
- [Configuring the IIS service group](#)
- [Configuring the MSVirtual Machine service group](#)
- [Configuring the service group for any additional applications](#)
- [Configuring an Oracle service group](#)

Configuring the File Share service group

Configuring the File Share service group involves creating a FileShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
 - Verify that the directories to be shared reside on shared drives.
 - Mount the drives containing the shared directories on the system where you run the wizard. Unmount the drives from other systems in the cluster.
 - Verify that Veritas Command Server service is running on all systems in the cluster.
 - Verify that you have the following information ready. The wizard will prompt you for this information:
 - A unique virtual computer name to be assigned to the file share server. This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
 - A unique virtual IP address to be assigned to the file share server. This is the IP address that the clients can use to access the file share.

Note: Windows Server 2008 does not support accessing file shares using the virtual IP address.

- The list of directories to be shared.
The wizard enables you to add existing shares to the VCS configuration. However, you cannot add special shares (shares created by the operating system for administrative and system use). For example, you cannot add the shares ADMIN\$, print\$, IPC\$, and *DriveLetter\$* to the VCS configuration.

- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).

Creating the File Share service group

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

To configure a FileShare

- 1 Start the File Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.

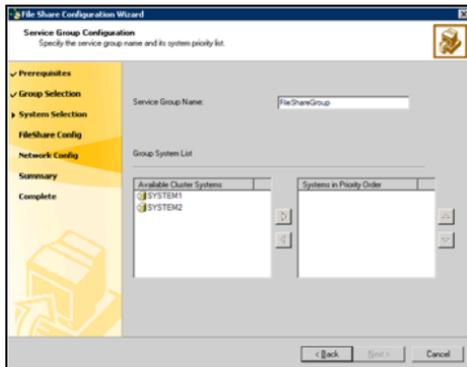
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > File Share Configuration Wizard**.

or

Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > File Share Configuration Wizard**.

- 2 Review the information in the Welcome panel and click **Next**.

- 3 On the Wizard Options panel, click **Create service group** and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the File Share service group.

Available Cluster Systems

Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

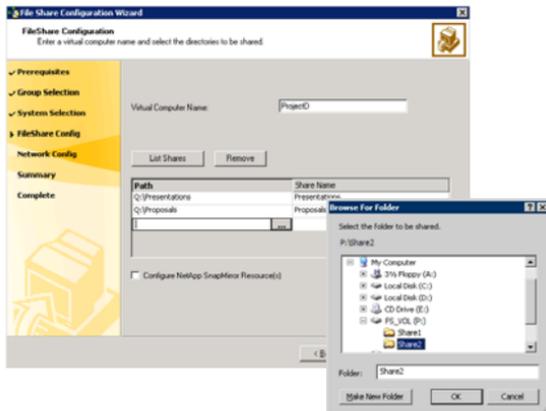
To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the File Share Configuration panel, specify the configuration information for the FileShare resources to be created and then click **Next**.

The wizard begins validating your configuration. Various messages indicate the validation status.

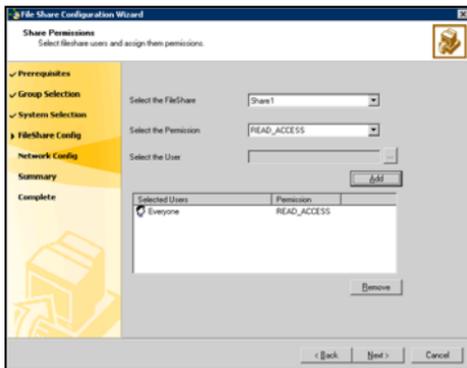


Virtual Computer Name Type a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 16 characters.

List Shares Click **List Shares** to view the existing shares on the shared storage, then select a share and click **Add**. You cannot add special shares (shares created by the operating system for administrative and system use).

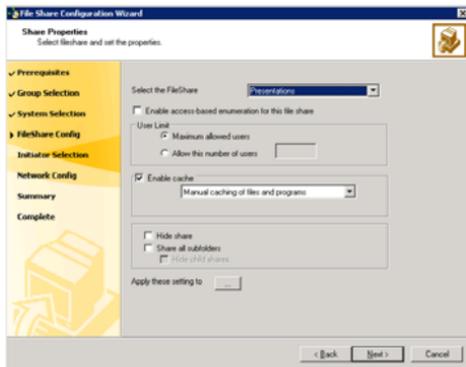
Path	<p>Click the field and either type the path of the directory to be shared or click the ellipsis button (...) to browse for a directory.</p> <p>The selected directories must meet the following conditions:</p> <ul style="list-style-type: none">■ The selected drive, the mount path, and the file path must not exist in the VCS configuration.■ The directories to be shared must reside on shared, non-system drives. <p>The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.</p>
Share Name	<p>If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can select a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.</p>
Remove	<p>To remove a file share from the configuration, click to select the file share and then click Remove.</p>

- On the Share Permissions panel, select the file share, specify the users for the file shares and assign permissions to them, and then click **Next**.



- | | |
|-----------------------|--|
| Select the FileShare | From the drop-down list, select the file share with which to associate user permissions, or select the default All FileShares to set the same permissions for all file shares. |
| Select the Permission | From the drop-down list, select the permission to be associated with the user. |
| Select the User | Click the ellipsis button (...), select a user, and click OK . |
| Add | Click Add to add the specified user to the Selected Users list. By default, all selected users are given READ_ACCESS permission. |
| Selected Users | Displays a list of selected users and the file share permissions. You can configure a maximum of 50 users for each file share. To configure more users, create a user group.
To change the file share permission associated with a user, click a user name in the Selected Users list and then select the desired permission from the Select the Permission drop-down list. |
| Remove | To deny file share access to a user, click the user name in the Selected Users list and click Remove . |

- 7 On the Share Properties panel, set the share properties for the file shares and then click **Next**.



Select the FileShare

From the drop-down list select a file share whose properties you wish to set.

Enable access-based enumeration for this file share

Check the **Enable access-based enumeration** check box to enable the Windows access-based enumeration feature on the selected file share.

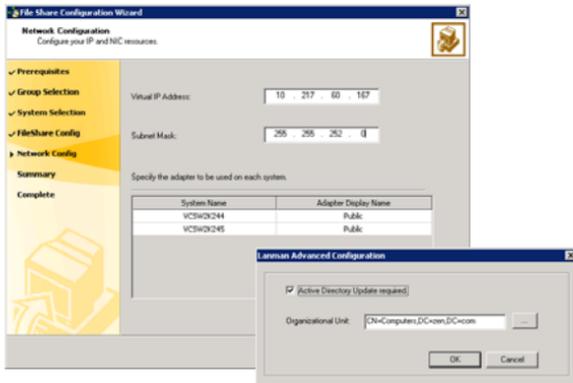
User Limit

Specify the number of users that are allowed access to the selected file share. Choose from the following options:

- **Maximum allowed users:** Select this option to allow access to the maximum numbers of users allowed on Windows.
- **Allow this number of users:** Select this option and then type the number of users that you wish to grant access to the selected file share. If you type zero or a value greater than what Windows supports, access is granted to the maximum allowed users on Windows.

Enable cache	<p>Check the Enable cache check box to enable local caching of the contents of the selected file share. Then, specify how the contents of the file share are available to users for offline access. In the drop down list select from the following caching options:</p> <ul style="list-style-type: none">■ Manual caching of files and programs: Only the files and programs specified by the user are available offline. This sets the FileShare resource attribute ClientCacheType to MANUAL.■ Automatic caching of programs: All the files and programs that the users access from the file share are available offline. This sets the FileShare resource attribute ClientCacheType to DOCS.■ Optimized automatic caching of files and programs: All files and programs, including executables, are cached locally. The next time the user accesses the executable files, they are launched from the local cache. This sets the FileShare resource attribute ClientCacheType to PROGRAMS.
Hide share	Check the Hide Share check box to make the new share a hidden share.
Share all subfolder	Check the Share all subfolders check box to share the subdirectories.
Hide child shares	Check the Hide child shares check box to hide the shared subdirectories.
Apply these settings to	<p>To apply the specified share properties to multiple file shares simultaneously, do the following:</p> <ol style="list-style-type: none">1 Click the ellipsis (...) button.2 On the Copy Share Properties dialog box, select the file shares from the Available Shares list and click the right arrow to move them to the Selected Shares list. Note that only those file shares that are not already shared are available for selection.3 Click OK. <p>Note: This option is not visible if you are configuring only one share in the service group.</p>

- 8 On the Network Configuration panel, specify information related to your network and then click **Next**.



Virtual IP Address

Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual IP belongs.

Adapter Display
Name

For each system in the cluster, select the public network adapter name.

This field displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. Verify that you select the adapters assigned to the public network, not the private.

Advanced Settings Click **Advanced Settings** to specify additional details for the Lanman resource.

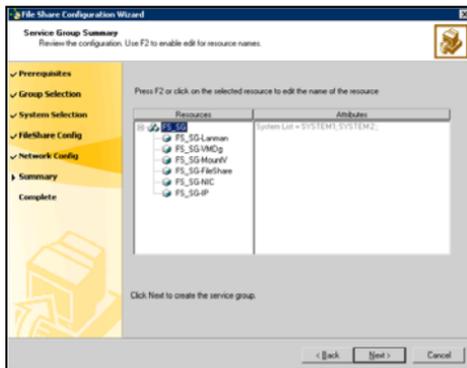
On the Lanman Advanced Configuration dialog box, complete the following:

- 1 Check **Active Directory Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
- 2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format
CN=containername,DC=domainname,DC=com.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
By default, the Lanman resource adds the virtual server to the default container "Computers."
- 3 Click **OK**.
The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- 9 On the Service Group Summary panel, review the service group configuration and click **Next**.

Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
- To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

- In the completion dialog box, check **Bring the service group online** check box if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the Print Share service group

Configuring the Print Share service group involves creating a PrintShare service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically
`C:\Program Files\Veritas\Cluster Server.`
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator* mode and then run the VCS commands.

Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Verify that VCS Command Server is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.
- Symantec recommends creating spooler and the replication directories on different disk partitions or volumes.
- Mount the drives with the spooler and the replication directories on the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).

- Verify that the software drivers for the network printers are installed on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - A unique virtual computer name to be assigned to the print share server.
This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
 - A unique virtual IP address to be assigned to the print share server.
 - The network printer's IP address.

Creating the Print Share service group

To create a Print Share service group perform the following tasks:

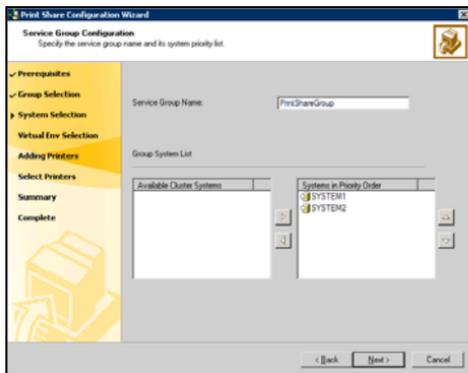
- Create a new service group with a PrintSpool resource and bring it online. This also involves configuring the Lanman resource on which the PrintSpool resource depends.
- Add a network printer to the virtual computer created by the Lanman resource. Create a new TCP/IP port for the printer.
- Configure a PrintShare resource in your service group and bring it online.

The following procedures describe these tasks in detail.

To create a Print Share service group with a PrintSpool resource

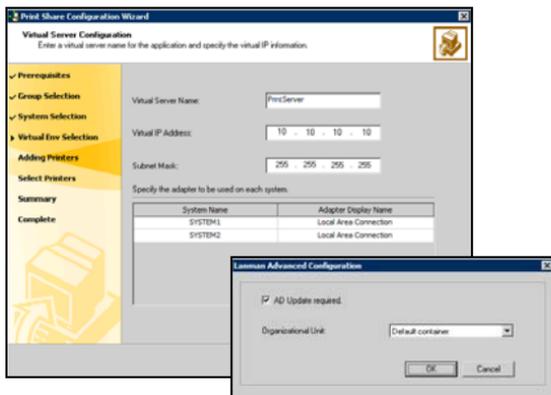
- 1 Start the Print Share Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > Print Share Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Print Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

- 4 On the Service Group Configuration panel, specify the service group details and click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name	Type a name for the Print Share service group.
Available Cluster Systems	Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list. To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow. To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Virtual Server Configuration panel, specify information related to your network and then click **Next**.



Virtual Server
Name

Type a unique virtual computer name by which the server will be known to clients. Note that the virtual name must not exceed 16 characters.

Virtual IP Address

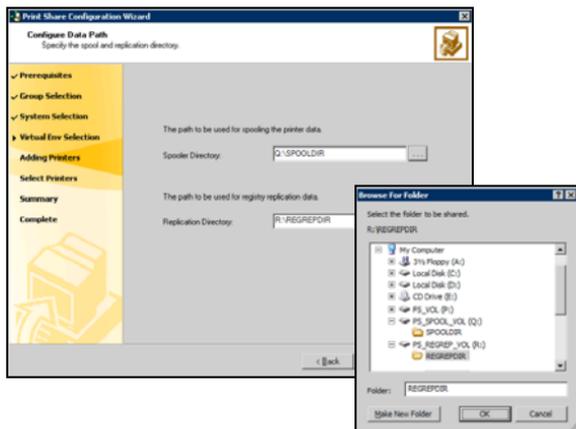
Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual server belongs.

Advanced Settings	<p>Click Advanced Settings to specify additional details for the Lanman resource.</p> <p>On the Lanman Advanced Configuration dialog box, complete the following:</p> <ol style="list-style-type: none">1 Check AD Update required check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format CN=containername,DC=domainname,DC=com. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."3 Click OK. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
Adapter Display Name	<p>Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow.</p> <p>For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.</p>

- 6 On the Configure Data Path panel, specify the spool and registry replication directories and then click **Next**.



Spooler Directory

Type the path or click ... (ellipsis button) to browse for the directory. All print commands will be spooled at this location.

Replication Directory

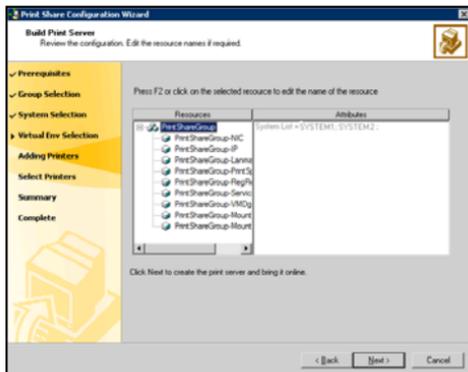
Type the path or click ... (ellipsis button) to browse for the directory. All changes related to the printer registry keys will be logged at this location.

The selected directories must fulfill the following conditions:

- The selected drive, the mount path, and the file path must not exist in the VCS configuration.
- The directories to be shared must reside on shared, non-system drives.
Symantec recommends creating the directories for replication and spooling on different mounts.

- 7 On the Build Print Server panel, review the configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**. The wizard starts running commands to add the PrintSpool resource and the resources on which the

PrintSpool resource depends, including the Lanman and ServiceMonitor resources.



- | | |
|------------|---|
| Resources | Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key. |
| Attributes | Displays the attributes and their configured values, for a resource selected in the Resources list. |

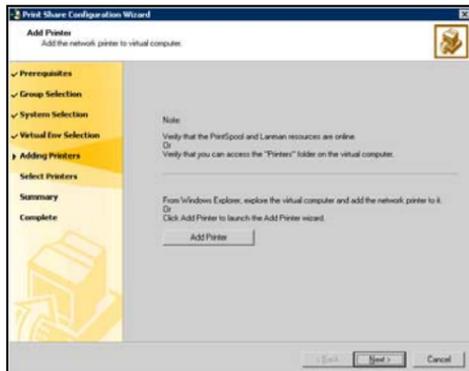
8 Bring the PrintSpool resource online.

Proceed to the next step to add the network printer to the virtual computer created by the Lanman resource and to create a new TCP/IP port for the printer.

To add the network printer to the virtual computer

- 1 Launch the Add Printer wizard to add the network printer to the virtual computer. Before starting the Add Printer wizard, verify that the PrintSpool and Lanman resources are online in your configuration.
To launch the Add Printer wizard, return to the Print Share Configuration Wizard and click **Add Printer** on the Add Printer panel, or in Windows Explorer, search for the virtual computer, explore the virtual computer by

double-clicking its name and on the virtual computer's Printers folder, double-click **Add Printer**.

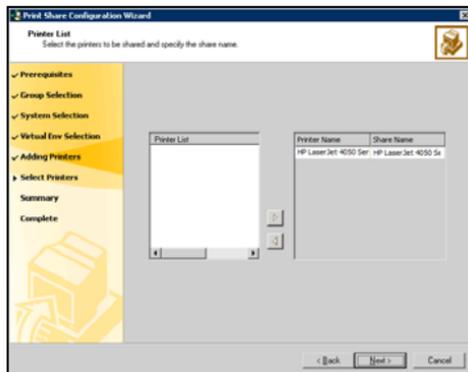


- 2 In the Add Printer wizard, review the information in the Welcome panel and click **Next**.
- 3 Follow the wizard instructions to add the network printer to the virtual computer.
In the Printer Sharing dialog box, always choose the **Do not share this printer** option.
Repeat these steps for each additional printer to be installed.
- 4 Return to the Print Share Configuration Wizard, and proceed to the next step to configure a PrintShare resource in your service group and bring it online.

To configure a PrintShare resource for the service group

- 1 On the Add Printer panel, click **Next**.

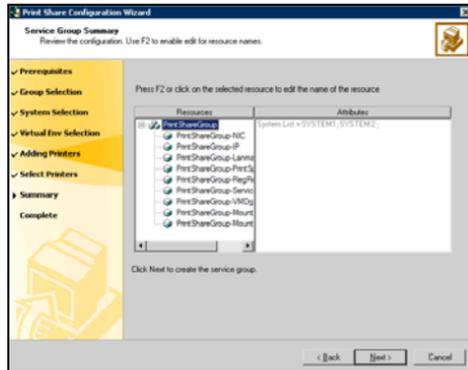
- 2 On the Printer List panel, specify the printers to be included in the Print Share service group and then click **Next**.



- Printer List** Click to select the printer, and then click the right arrow to include the selected printers in your service group. To remove a selected printer from your service group, click the printer from the Printer Name list and click the left arrow.
- Share Name** Type a unique share name for the printer by which it will be known to clients. If you previously chose to share the printer, VCS uses the printer's share name.

- 3 On the Service Group Summary panel, review the service group configuration and then click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration.

Click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



Resources

Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes

Displays the attributes and their configured values, for a resource selected in the Resources list.

- 4 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the IIS service group

Configuring the IIS service group involves creating a IIS service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify the sites to be monitored are on shared storage.
- For IIS 7.0 on Windows Server 2008, you must install the following components:

- IIS 6 Metabase Compatibility
- IIS 6 WMI Compatibility

These options are available under Management Tools on the Role Services page of the Add Roles Wizard.

These components are required for the IIS agent to function on Windows Server 2008.

For IIS 7.0 on Windows Server 2008 Server Core, you must install IIS in the specified order. See ["Installing IIS 7.0 on Windows Server 2008 Server Core"](#) on page 286 for instructions.

- Do not use the IIS agent to configure SMTP and NNTP sites if you have Microsoft Exchange installed.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- Verify the port numbers assigned to IIS sites are not used by other programs.
- Synchronize the IIS configuration on all nodes hosting the service group. See ["Synchronizing IIS configuration on Windows 2003"](#) on page 285 for instructions.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- To configure IIS agent on Windows Server 2008 Server Core, you must manually add the required resources and configure the service group. You can perform the manual configuration steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Mount the drives containing the shared directories from the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - IIS sites to be monitored.
 - Application pools associated with each site.
 - Port numbers associated with each site.
 - Virtual IP addresses and computer names associated with the sites. The virtual IP addresses and the virtual computer names must have forward and reverse entries in the DNS.
- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Synchronizing IIS configuration on Windows 2003

Complete the following steps.

To synchronize the IIS configuration on Windows 2003 systems

Synchronize the IIS configuration on all nodes that will host the IIS service group.

- 1 Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system. For example, the following command copies the IIS metabase to *target_system*. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing IIS 7.0 on Windows Server 2008 Server Core

Complete the following steps.

To install IIS 7.0 on Windows Server 2008 Server Core

- 1 Type the following at the command prompt:

```
start /w pkgmgr /iu:IIS-WebServerRole;IIS-WebServer;IIS-  
CommonHttpFeatures;  
IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;  
IIS-HttpErrors;IIS-HttpRedirect;IIS-ApplicationDevelopment;  
IIS-ASP;IIS-CGI;IIS-ISAPIExtensions;IIS-ISAPIFilter;  
IIS-ServerSideIncludes;IIS-HealthAndDiagnostics;  
IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;  
IIS-HttpTracing;IIS-CustomLogging;IIS-ODBCLogging;IIS-  
Security;IIS-BasicAuthentication;IIS-WindowsAuthentication;  
IIS-DigestAuthentication;  
IIS-ClientCertificateMappingAuthentication;  
IIS-IISCertificateMappingAuthentication;  
IIS-URLAuthorization;IIS-RequestFiltering;IIS-IPSecurity;  
IIS-Performance;IIS-HttpCompressionStatic;  
IIS-HttpCompressionDynamic;IIS-WebServerManagementTools;  
IIS-ManagementScriptingTools;IIS-IIS6ManagementCompatibility;  
IIS-Metabase;IIS-WMICompatibility;IIS-LegacyScripts;  
IIS-FTPPublishingService;WAS-WindowsActivationService;  
IIS-FTPPublishingService;IIS-FTPService
```

- 2 Verify that all the components specified in the earlier step have successfully installed. Type the following at the command prompt:

```
notepad C:\windows\logs\cbs\cbd.log
```

This opens the log file, `cbd.log`, in the Notepad text editor.

- 3 Check the entries in the log file, `cbd.log`. The last log entry should resemble the following:

```
Info CBS Pkgmgr: return code: 0x0
```

This message indicates that all the components are installed successfully.

- 4 Run the `oclist` command to verify that the following components are installed:
IIS-WebServerRole; IIS-WebServer; IIS-IIS6ManagementCompatibility; IIS-Metabase; IIS-WMICompatibility; IIS-FTTPublishingService; WAS-WindowsActivationService; IIS-FTTPublishingService; IIS-FTPService
Type the following at the command prompt:

```
oclist
```
- 5 Repeat the steps on all the nodes on which you wish to configure the IIS service group.

Creating the IIS service group

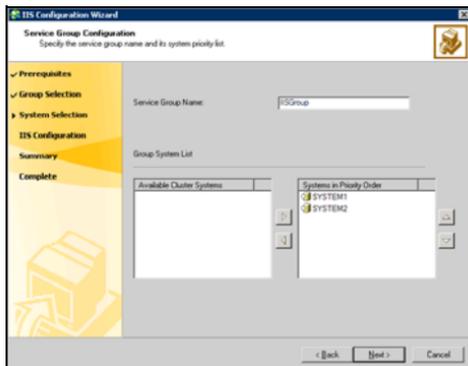
Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for information on resource types, attribute definitions, resource dependencies, and sample service group configurations.

The following steps describe how to create an IIS service group.

To create an IIS service group

- 1 Start the IIS Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > IIS Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > IIS Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the IIS service group.

Available Cluster Systems

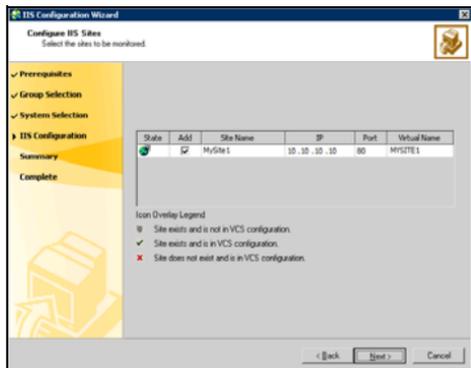
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

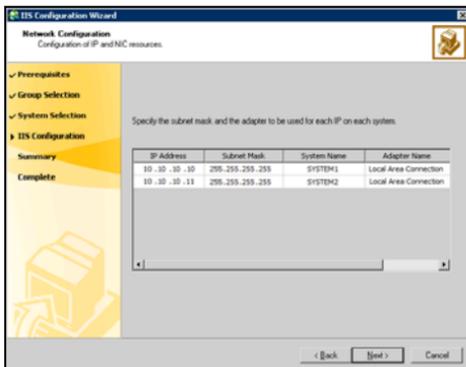
System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names, and then click **Next**.



- Add** Check the check box corresponding to the site to be configured in VCS.
- IP** Type the virtual IP address for each site to be configured. Make sure that each virtual IP address is associated with only one virtual computer name and vice-versa.
- Port** Type the port number for each site to be configured.

- 6 On the Network Configuration panel, specify information related to the virtual IP addresses and click **Next**.



IP Address

Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems.

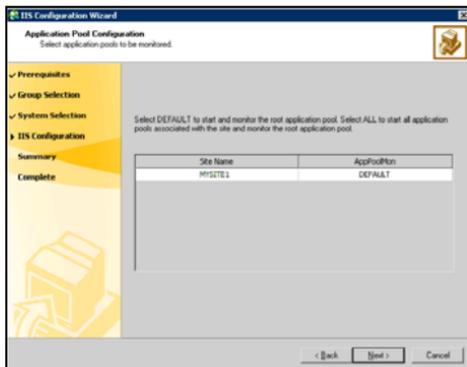
Subnet Mask

Type the subnet mask associated with each virtual IP address.

Adapter Name

Select the adapter associated with the virtual IP address on each system.

- 7 On the Application Pool Configuration panel, select the monitoring options for application pools associated with each site and click **Next**.



Site Name

Displays the site names.

AppPoolMon

For each site, select the monitoring options from the AppPoolMon list.

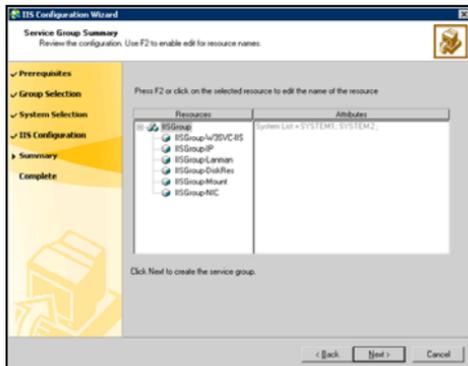
NONE—The agent will not monitor the application pool associated with the site.

DEFAULT—Starts and monitors the root application pool associated with the site.

ALL—Starts all application pools associated with the site and monitors root application pool.

- 8 On the Service Group Summary panel, review the service group configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click

Yes. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



Resources

Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes

Displays the attributes and their configured values, for a resource selected in the Resources list.

- 9 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the MSVirtual Machine service group

Configuring the MSVirtual Machine service group involves creating a MSVirtual Machine service group and defining the attribute values for its resources. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141
For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- Verify that the shared drives required by the applications are mounted.
- Mount the drives containing the shared directories of the virtual machine, on which the wizard will be run. Unmount the drives from other systems in the cluster.

- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
- Disable the firewall on each node that will host the service group.
- You must have the following information ready. The wizard will prompt you for this information:
 - The name of the virtual machine.
 - Destination on shared disks for the virtual hard disk files.
 - Network adapters on physical nodes to be associated with network adapters on the virtual machine.
 - Information about monitoring heartbeats (optional).

Creating the MSVirtual Machine service group

The following steps describe how to create an IIS service group.

To create the MSVirtualMachine service group

- 1 Start the MSVirtual Machine Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > MSVirtual Machine Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > MSVirtual Machine Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, select the **Create service group** option and click **Next**.
- 4 Enter a name for the service group and specify the systems on which to configure the service group.
 - Enter a name for the service group.
 - In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the **Systems in Priority Order** box and click the left arrow.

- To change a system's priority in the service group's system list, click the system from the **Systems in Priority Order** and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.
- 5 Specify details about the virtual machine.
- Select the virtual machine.
 - For each virtual disk, specify a destination folder where the virtual hard disk files will be moved. Click the Browse icon to browse for folders.
 - To enable detail monitoring for the virtual machine, select the **Monitor Heartbeats** check box and enter failed heartbeat threshold in the **No. of Monitor Cycles** field.
The threshold defines the number of consecutive monitor cycles the agent waits to detect heartbeats from the virtual machine before declaring the resource as faulted.
 - Click **Next**.
- 6 Select an adapters corresponding to the virtual machine on each system.
- For each system in the cluster, enter or click a network adapter name to be associated with the network adapters on the virtual machine.
To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.
The fields for the virtual IP address and subnet mask are disabled by design.
 - Click **Next**.
- 7 Review the service group configuration.
The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.
- The wizard assigns unique names to resources. Change names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.
 - Click **Next**.

- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.
The wizard starts running commands to create the service group. Various messages indicate the status of these commands.
- 8 In the completion dialog box, select the check box if you want to bring the service group online on the local system.
- 9 Click **Finish**.

Configuring the service group for any additional applications

Configuring the service group for any additional application involves creating an application service group and defining the attribute values for its resources. This can be done using the Application Configuration Wizard. After the service group is created, you must configure the shares to mount automatically at startup.

Prerequisites

- The application is installed on all the nodes that are going to be part of the service group.
 - The shared drives required for the application are mounted on this node.
 - The startup type of the application service is set to manual on all nodes in the service group.
 - The application service is stopped on all nodes in the service group.
 - If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the

steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console).

Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Before running the wizard, make sure you have the following information ready:
 - Details (service name, start parameters, startup directory, etc.) of the application that you wish to configure.
 - Shared storage used by the applications.
 - Application registry entries for configuring registry replication.
 - Network and virtual computer (Lanman) details for the application.

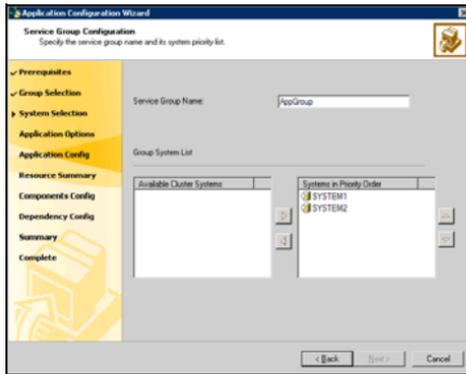
Creating the application service group

The following steps describe how to create an IIS service group.

To create an application service group

- 1 Start the Application Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > Application Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Application Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Create service group** and click **Next**.

4 Specify the service group name and system list.

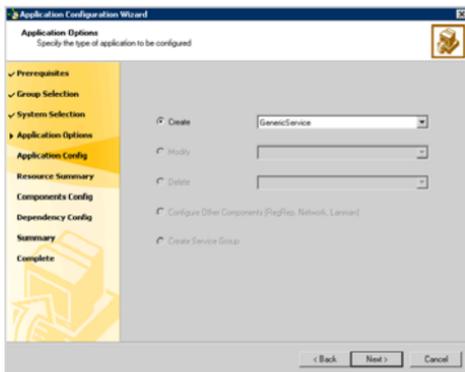


- Enter a name for the service group.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.
To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
 - To change a system's priority in the service group's system list, select the system in the Systems in Priority Order list and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 The Application Options panel provides you the option to specify the type of application to be configured. The available options are:
- **Generic Service:** Configures a service using the Generic Service agent. The agent brings services online, takes them offline, and monitors their status. See "[Configuring a GenericService resource](#)" on page 299.
 - **Process:** Configures a process using the Process agent. The agent brings processes online, takes them offline, and monitors their status. See "[Configuring processes](#)" on page 303.
 - **Service Monitor:** Configures a service using the ServiceMonitor agent. The agent monitors a service or starts a user-defined script and

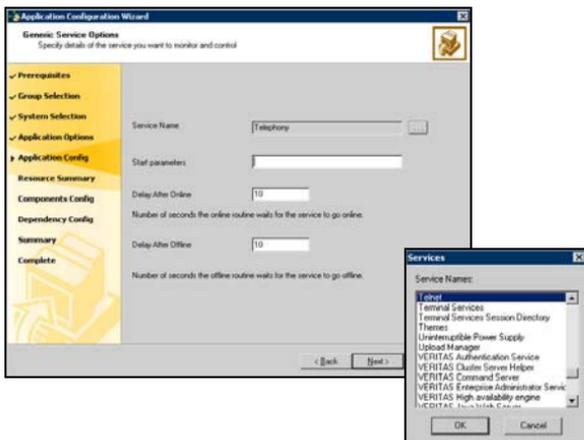
interprets the exit code of the script. See “[Configuring a ServiceMonitor resource](#)” on page 307.

Configuring a GenericService resource

- 1 In the Application Options panel, click **Create**, select **GenericService** from the corresponding drop-down list, and click **Next**.

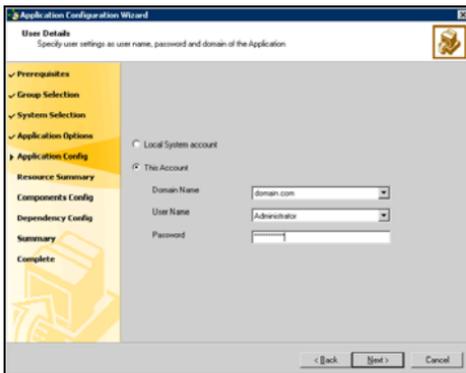


- 2 Select the service name for which you wish to configure a GenericService resource. Also specify the attributes for the resource.



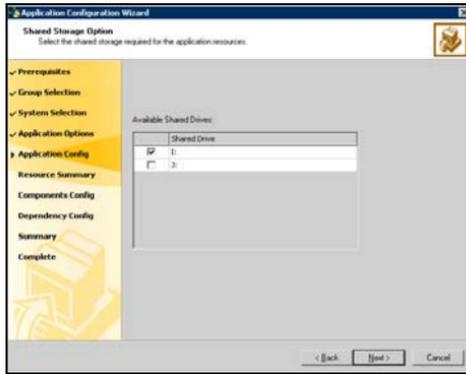
- Click the icon (...) adjacent to the Service Name text box.
- In the Services dialog box, select a service and click **OK**. The selected service appears in the Service Name text box.
- In the **Start Parameters** text box, provide the start parameters for the service, if any.
- In the **Delay After Online** text box, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
- In the **Delay After Offline** text box, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
- Click **Next**.

- 3 Specify the information about the user in whose context the service will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the GenericService resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

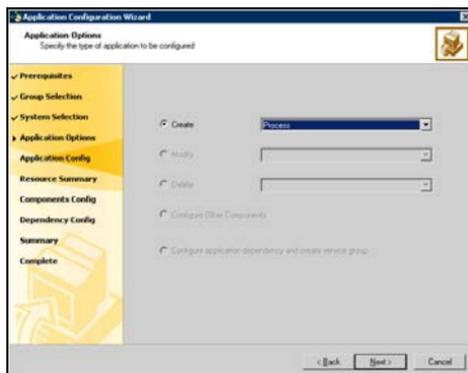


- In the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another GenericService resource, repeat [step 1](#) through [step 5](#).
 - To configure a Process resource, proceed to “[Configuring processes](#)” on page 303 for instructions.
 - To configure a ServiceMonitor resource, proceed to “[Configuring a ServiceMonitor resource](#)” on page 307 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 310 for instructions.

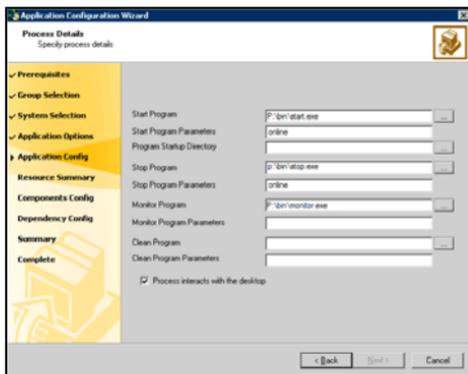
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 313.

Configuring processes

- 1 In the Application Options panel, click **Create**, select **Process** from the corresponding list, and click **Next**.



- 2 Specify the details for the process.



- In the **Start Program** text box, specify the complete path of the program that will start the process to be monitored by VCS. You can

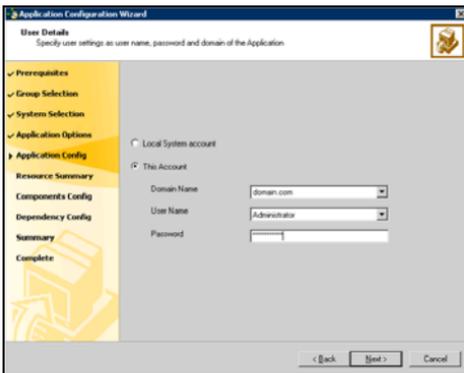
choose to either type in the location of the program or browse for it using the (...) icon.

- In the **Start Program Parameters** text box, specify the parameters used by the Process agent start program.
- In the **Program Startup Directory** text box, enter the complete path of the Process agent program or browse for it by clicking the (...) icon.
- In the **Stop Program** text box, enter the complete path of the program that will stop the process started by the Start Program or browse for it by clicking the (...) icon.
- In the **Stop Program Parameters** text box, specify the parameters used by the stop program.
- In the **Monitor Program** text box, enter the complete path of the program that monitors the Start Program or browse for it by clicking the (...) icon.

If you do not specify a value for this attribute, VCS monitors the Start Program. If the Start Program is a script to launch another program, you must specify a monitor program.

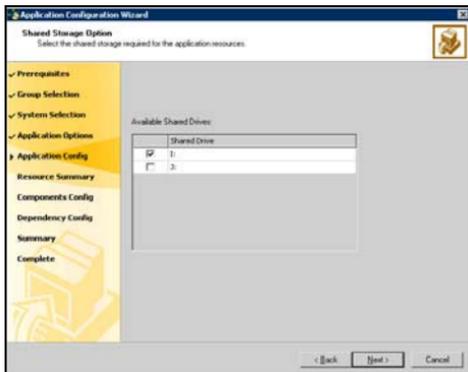
- In the **Monitor Program Parameters** text box, specify the parameters used by the monitor program.
- In the **Clean Program** text box, enter the complete path of the Clean process or browse for it by clicking the (...) icon.
- If no value is specified, the agent kills the process indicated by the Start Program.
- In the **Clean Program Parameters** text box, specify the parameters used by the Clean program.
- Select the **Process interacts with the desktop** check box if you want the process to interact with your Windows desktop. Setting this option enables user intervention for the process.
- Click **Next**.

3 Specify information about the user in whose context the process will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the Process resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

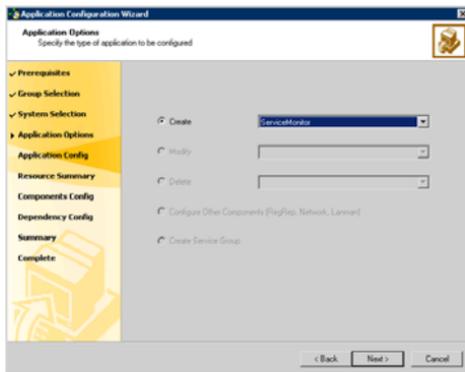


- From the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the Process resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another Process resource, repeat [step 1](#) through [step 5](#).
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 299 for instructions.
 - To configure a ServiceMonitor resource, proceed to [“Configuring a ServiceMonitor resource”](#) on page 307 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to [“Configuring VCS components”](#) on page 310 for instructions.

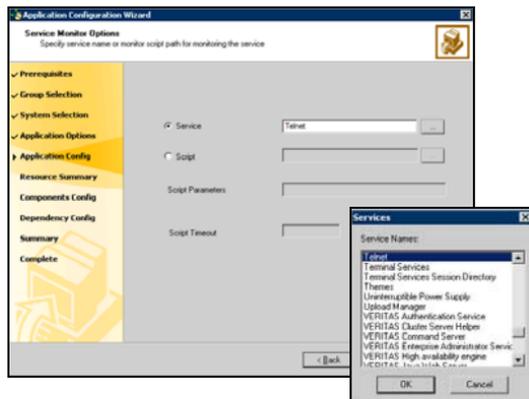
If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 313.

Configuring a ServiceMonitor resource

- 1 In the Application Options panel, click **Create**, select **ServiceMonitor** from the corresponding drop-down list, and click **Next**.



- 2 Specify the service to be monitored or a user-defined script to monitor a service.



If you want VCS to monitor the service:

- Select the **Service** option and click the icon (...) adjacent to the **Service Name** text box.
- In the Service dialog box, select the service and click **OK**. The selected service name appears in the **Service Name** text box. Alternatively, You may also type in the service name to be monitored.
- Click **Next**.

If you want a script to monitor the service:

- Specify the complete path for the script using the Browse button (...).
- Specify the parameters for the script.
- Specify the time in seconds for the agent to receive a return value from the monitor script.
- Click **Next**.

- 3 Specify the user information in whose context the service will be monitored.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.

If the service selected in [step 2](#) on page 307 is running in the context of a local system account, the **This Account** option is disabled. Similarly, if the service is running in the context of any other user account, the **Local System account** option is disabled.

- Click **Next**.

ServiceMonitor resource belongs to the category of *persistence* resources. Such resources do not depend on other VCS resources, including shared storage. Hence, the Shared Storage Option panel does not appear if you select the ServiceMonitor option.

- 4 In the Application Resource Summary panel, review the summary of the ServiceMonitor resource. Click **Back** to make changes. Otherwise, click **Next**.
- 5 The Application Options panel appears. Select one of the following options:
 - To configure another ServiceMonitor resource, repeat [step 1](#) through [step 4](#).
 - To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 299 for instructions.
 - To configure a Process resource, see “[Configuring processes](#)” on page 303 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 310 for instructions.

If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 313.

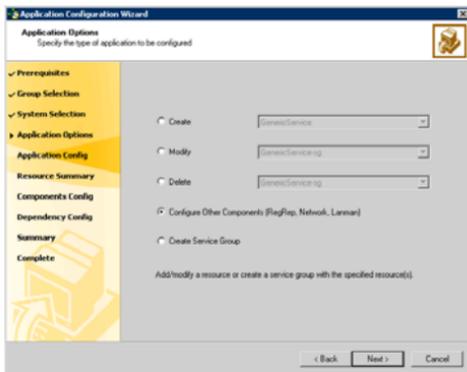
Configuring VCS components

Applications configured using GenericService or Process resources may require network components, or Registry Replication resources. You can configure these VCS components *only* for service groups created using the wizard.

Note: Configure these components only after configuring all application resources. The wizard creates a service group after these components are configured. To add more application resources, you must rerun the wizard in the Modify mode.

To configure VCS components

- 1 In the Application Options panel, click **Configure Other Components**.



- 2 Select the VCS component to be configured for your applications.

The available options are:

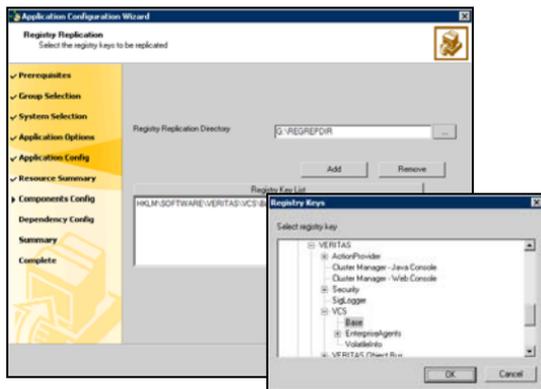
- **Registry Replication Component:** Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to [step 3](#) on page 311.
- **Network Component:** Select this option to configure network components for your application. If you wish to configure a virtual computer name, check **Lanman component** also. To configure a network resource, proceed to [step 4](#) on page 312.

The wizard does not enable the **Lanman Component** check box unless the **Network Component** check box is checked.

To configure Registry Replication

The RegistryReplication panel appears only if you chose to configure the Registry Replication Component in the Application Component panel.

- 3 Specify the registry keys to be replicated.



- Specify the directory on the shared disk in which the registry changes are logged.
- Click **Add**.
- In the Registry Keys dialog box, select the registry key to be replicated.
- Click **OK**.
- The selected registry key is added to Registry KeyList box. Click **Next**.

If you chose Network Component from the Application Component panel, proceed to the next step. Otherwise, proceed to [step 5](#) on page 312.

To configure network components

The Virtual Computer Configuration panel appears only if you chose to configure the Network Component in the Application Component panel.

- 4 Specify information related to your network.
 - Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 16 characters.

Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component panel.
 - Enter a unique virtual IP address for the virtual server.
 - Enter the subnet to which the virtual server belongs.
 - Click **Advanced** to specify additional details for the Lanman resource.
 - Check **AD Update required** to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
 - In the Organizational Unit field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
 - For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

Note that the wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private.
 - Click **Next**.
- 5 The Application Options panel is displayed. Select one of the following options:

- To configure additional VCS components, repeat [step 1](#) on page 310 through [step 4](#) on page 312.
- To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 299 for instructions.
- To configure a Process resource, see “[Configuring processes](#)” on page 303 for instructions.
- To configure a Service Monitor resource, see “[Configuring a ServiceMonitor resource](#)” on page 307 for instructions.

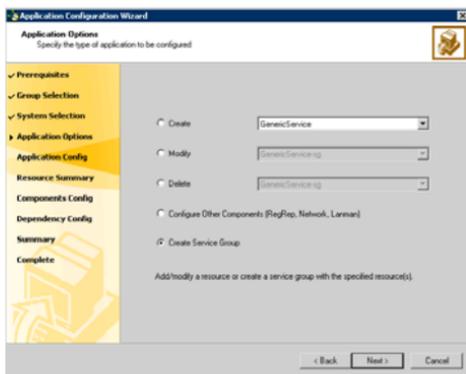
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 313.

Configuring Application Dependencies

The Application Configuration Wizard enables you to create service group for the application resources and other VCS components configured using the wizard. This section describes how to create the service group using the wizard.

To create a service group

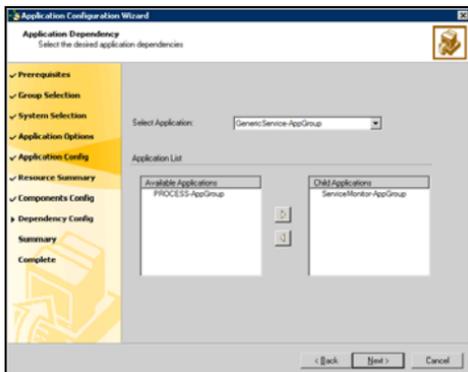
- 1 In the Application Options panel, click **Configure application dependency and create service group**.



The option is enabled only if:

- resources and VCS components are already configured using the wizard.

- you clicked **Modify Service Groups** in the Wizard Options panel.
- 2 Specify the dependency between the applications.

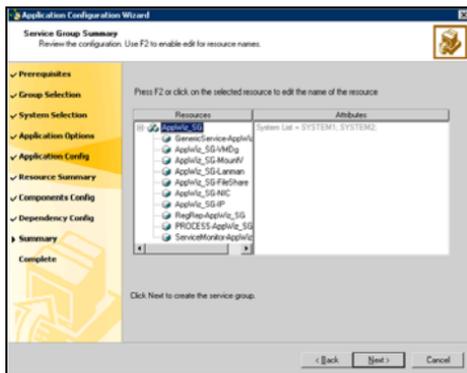


You must have at least two resources configured for Application Dependency panel to appear. Of the two resources, one should either be a GenericService or a Process resource.

- From the Select Application list, select the application that would depend on other applications. The selected application becomes the parent application.
- From the Available Applications list, select the application on which the parent application would depend and click the right-arrow icon to move the application to the Child Applications list.
- To remove an application from the Child Applications list, select the application in the list and click the left arrow.
- Repeat these steps for all such applications for which you want to create a dependency.
- Click **Next**.

The Application Dependency panel enables you to link resources configured using the wizard. If these resources are dependent on other services outside the VCS environment, you should first configure resources for such services and then create the appropriate dependency.

3 Review the service group configuration.



The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

- The wizard assigns unique names to resources. Change names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.
 - Click **Next**.
 - A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.
The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 4 In the completion panel, check **Bring the service group online** if you want to bring the service group online on the local system.
 - 5 Click **Finish** to create the service group and exit the Application Configuration Wizard.

Configuring an Oracle service group

Configuring the Oracle database agent involves creating the Oracle service group, its resources, and defining attribute values for the configured resources.

VCS provides several ways to configure the Oracle agent: the agent configuration wizard, the Java and Web consoles, and the command line. This section provides instructions on how to use the agent configuration wizard to configure the agent.

Prerequisites

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
 - You must be a Domain Administrator on the node where you run the wizard.
 - If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Verify that SFW HA, along with the VCS database agent for Oracle, is installed on all cluster nodes.
 - Verify a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
 - Verify that the Veritas high availability engine (HAD) is running on the system from where you run the wizard.
 - Mount the shared drives containing the data files, control files, redo log files, bdump, cdump, and udump files. Unmount the drives from other systems in the cluster.
 - If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
 - Assign the virtual IP address on the system where you run the wizard. Remove the virtual IP address from other systems

- Mount the database and start the Oracle instance on the node running the wizard.
- Make sure that you have the following information ready; the wizard will prompt you for this information:
 - The databases and listeners to be monitored by VCS.
 - For the instances to be monitored in detail, name and location of the respective SQL files.
 - A valid domain name, user name, and password with which the database service was configured for the database.

Creating an Oracle service group

This following procedure describes how to create an Oracle service group.

To create an Oracle service group

- 1 Start the Oracle Configuration wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
From the **Solutions Configurations Center** expand the Solutions for Additional Applications and from the display click **Campus Cluster Configuration > Configure the Service Group > Oracle Agent Configuration Wizard**.
or
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Oracle Agent Configuration Wizard**.
- 2 In the Welcome panel click **Next**.
- 3 In the Wizard Options panel, select **Create service group** and click **Next**.
- 4 In the Service Group Configuration panel, complete the following and click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

Service Group Name

Type a name for the Oracle service group.

Available Cluster Systems Select the systems on which to configure the service group and click the right arrow to move the systems to the Systems in Priority Order box.

The Systems in Priority Order box represents the service group's system list. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

To remove a system from the service group's system list, select a system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, select the system from the Systems in Priority Order box, and click the up and down arrows.

- 5 In the Oracle Configuration panel, select the SIDs and the associated listeners to be added to the service group, and click **Next**.
The SID is a system identifier that uniquely identifies the Oracle database instance, while the listener is the name of the corresponding listener service.
- 6 On the Detail Monitoring panel, configure detail monitoring for the Oracle database if required, and click **Next**.

Detail Monitor Check the Detail Monitor option corresponding to each database that you want to configure detail monitoring for.

SQL Path Type the path of the SQL file that will query the database to validate the status. Click the icon next to the field to browse for the SQL file.

A sample SQL file, check.sql, is located at
`%VCS_HOME%\bin\Oracle\.`

- 7 In the Domain and User selection panel, type a valid domain name, user name, and password with which the database service was configured for the database and click **Next**.
- 8 On the Network Configuration panel, specify the network related information and click **Next**.

The wizard discovers and displays the virtual IP address for the Oracle Server.

Subnet Mask	Type the subnet mask to which the virtual IP belongs.
Adapter Display Name	For each system in the cluster, select the public network adapter. Select the Adapter Name field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- 9 Review the configuration on the Summary panel and click **Next**.

Resources	Lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box. The wizard assigns unique names to the resources. To edit a resource name, select the resource name and click on it, or press the F2 key. After the edit, press the Esc key to cancel the changes, or press the Enter key to confirm the changes.
-----------	---

- 10 In the confirmation dialog box, click **Yes**. Click **No** if you wish to review your settings.
The wizard starts running commands to create the Oracle service group.
- 11 On the Completing the Oracle Configuration panel, check **Bring the service group online** to bring the service group online on the local system, and click **Finish**. The Oracle service group is created in your cluster.

Configuring dependent services

If the database service has other dependent services, make sure the dependent services are running on the node where the database service is online. Note that the online agent operation brings only the database service online and not the dependent services.

For example, on Oracle 10g, the DBConsole service corresponding to an Oracle database has a dependency on the database service. That is, for the Enterprise Manager to manage the databases, you must make sure the DBConsole service is running on the node where the database service is online.

To configure a dependent service

- 1 For the dependent service, add a GenericService resource manually.

- 2 Make the GenericService resource dependent on the corresponding Oracle resource.
- 3 Set the Critical attribute to False if the Oracle service group must not fail over when the GenericService resource faults.

Refer to the *Veritas Cluster Server Administrator's Guide* for configuration instructions.

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 Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in [step 1](#).
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.
The service group you selected is taken offline and brought online on the node that you selected.

Replicated Data Clusters

This section includes the following chapters:

- [About Replicated Data Clusters](#)
- [Deploying Replicated Data Clusters: New application installation](#)

About Replicated Data Clusters

This chapter includes the following topics:

- [“About Replicated Data Clusters”](#) on page 326
- [“How VCS Replicated Data Clusters work”](#) on page 327
- [“Setting up a Replicated Data Cluster configuration”](#) on page 328
- [“Migrating the service group”](#) on page 331

About Replicated Data Clusters

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage, to assure data access to all the nodes in a cluster.

The Replicated Data Cluster configuration provides both local high availability and disaster recovery functionality in a single VCS cluster. You can set up RDC in a VCS environment using Veritas Volume Replicator (VVR).

An RDC exists within a single VCS cluster with a primary zone and a secondary zone, which can stretch over two buildings or data centers connected with Ethernet. In an RDC configuration, if an application or a system fails, the application is failed over to another system within the current primary zone. If the entire primary zone fails, the application is migrated to a system in the secondary zone (which then becomes the new primary).

For VVR replication to occur, the disk groups containing the Replicated Volume Group (RVG) must be imported at the primary and secondary zones. The replication service group must be online at both zones simultaneously, and must be configured as a hybrid VCS service group.

The application service group is configured as a failover service group. You must configure the application service group with an online local hard dependency on the replication service group.

Note: VVR supports multiple replication secondary targets for any given primary. However, RDC for VCS supports only one replication secondary for a primary.

An RDC configuration is appropriate in situations where dual dedicated LLT links are available between the primary zone and the secondary zone but lack shared storage or SAN interconnect between the primary and secondary data centers. In an RDC, data replication technology provides node access to data in a remote zone.

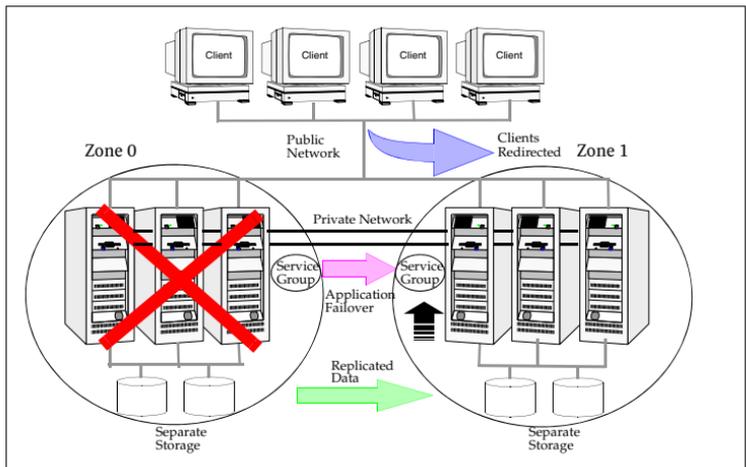
You must use dual dedicated LLT links between the replicated nodes.

How VCS Replicated Data Clusters work

To understand how a RDC configuration works, let us look at an application configured in a VCS replicated data cluster. The configuration has two system zones:

- Primary zone (zone 0) comprising nodes located at the primary site and attached to the primary storage
- Secondary zone (zone 1) comprising nodes located at the secondary site and attached to the secondary storage

The application is installed and configured on all nodes in the cluster. The application data is located on shared disks within each RDC zone and is replicated across RDC zones to ensure data concurrency. The application service group is online on a system in the current primary zone and is configured to fail over in the cluster.



If the system or application fails, VCS attempts to fail over the application service group to another system within the same RDC system zone. However, if VCS cannot find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone (zone 1). VCS also redirects clients once the application is online on the new location.

Setting up a Replicated Data Cluster configuration

In the example, the application is configured as a VCS service group in a four-node cluster, with two nodes in the primary RDC zone and two in the secondary RDC zone. If a failure occurs on the primary node, VCS can fail over the application to the second node in the primary zone.

The process involves the following tasks:

- [Setting up replication](#)
- [Configuring the service groups](#)

Setting up replication

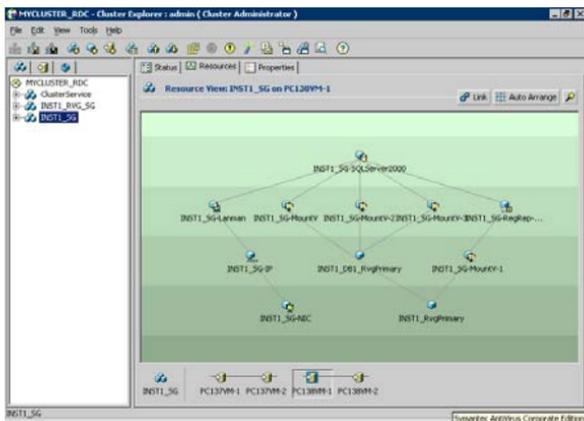
Set up replication between the shared disk groups. Use VVR to group the shared data volumes into a Replicated Volume Group, and creating the VVR Secondary on hosts in your secondary zone.

Create a Replicated Data Set (RDS) with the Primary RVG consisting of the shared volumes between the nodes in the first zone and Secondary RVG consisting of shared volumes between nodes in the second zone. Therefore, use the same Disk Group and RVG name in both zones so that the MountV resources will mount the same block devices.

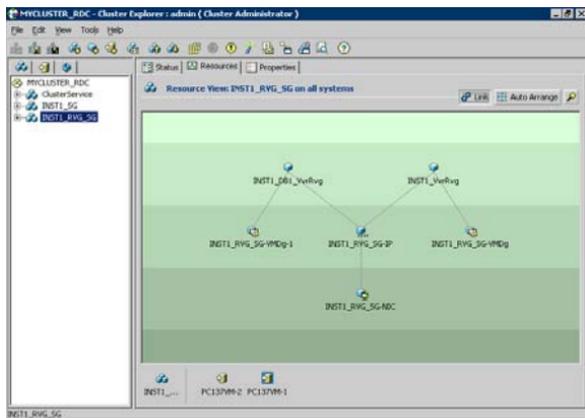
Configuring the service groups

For a successful wide-area failover, the mount points and applications must fail over to the secondary RDC zone. Additionally, the VVR secondary disk group and RVG must be imported and started on the secondary RDC zone.

The following screen from the VCS Cluster Manager (Java Console) depicts a typical application service group RDC configuration. This example uses the SQL Server application; however, the basic concepts are same, regardless of the application.



The following screen from the VCS Cluster Manager (Java Console) depicts a typical replication service group (RVG) configuration, again using SQL Server as an example:



Migrating the service group

In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the shared storage. In this situation, none of the nodes in the primary zone see any device.

The service group cannot fail over locally within the primary RDC zone, because the shared volumes cannot be mounted on any node. So, the service group must fail over to a node in the current secondary RDC zone.

The RVGPrimary agent ensures that VVR volumes are made writable. The application can be started at the secondary RDC zone and run there until the problem with the local storage is corrected.

If the storage problem is corrected, you can switch the application back to the primary zone using VCS.

Before you switch the application back to the original primary RDC zone, you must resynchronize any changed data from the active secondary RDC zone since the failover. Once the resynchronization completes, switch the service group to the primary zone.

To switch the service group

- 1 In the **Service Groups** tab of the of the Cluster Explorer configuration tree, right-click the service group.
- 2 Click **Switch To** and select the system in the primary RDC zone to switch to.
- 3 Click **OK**.

Deploying Replicated Data Clusters: New application installation

This chapter includes the following topics:

- [“Tasks for a new replicated data cluster installation—additional applications”](#) on page 334
- [“Reviewing the requirements”](#) on page 336
- [“Reviewing the configuration”](#) on page 342
- [“Configuring the storage hardware and network”](#) on page 343
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 346
- [“Configuring VxSAS”](#) on page 352
- [“Configuring the cluster”](#) on page 355
- [“Configuring disk groups and volumes”](#) on page 370
- [“Installing and configuring the application or server role”](#) on page 379
- [“Configuring the service group”](#) on page 382
- [“Creating the primary system zone”](#) on page 433
- [“Verifying the installation in the primary zone”](#) on page 434
- [“Creating a parallel environment in the secondary zone”](#) on page 435
- [“Adding the systems in the secondary zone to the cluster”](#) on page 436
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 441
- [“Configuring a hybrid RVG service group for replication”](#) on page 452

- [“Setting a dependency between the service groups”](#) on page 462
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 463
- [“Verifying the RDC configuration”](#) on page 469
- [“Additional instructions for GCO disaster recovery”](#) on page 470

Tasks for a new replicated data cluster installation—additional applications

Configure the high availability and application components on the primary and secondary zones, then complete the Replicated Data Set solution by configuring the components for both zones.

For more information on VVR, see the *Veritas Volume Replicator Administrator’s Guide*.

[Table 12-1](#) outlines the high-level objectives and the tasks to complete each objective:

Table 12-1 Task List: New RDC configuration

Objective	Tasks
“Reviewing the requirements” on page 336	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 342	<ul style="list-style-type: none"> ■ Understanding active/passive configuration and zone failover in a RDC environment ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 343	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which the application will be installed
“Installing Veritas Storage Foundation HA for Windows” on page 346	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ■ If applicable, selecting the option to install the Veritas Cluster Server Agent for your application
“Configuring VxSAS” on page 352	<ul style="list-style-type: none"> ■ Using the VVR Security Service Configuration wizard to configure the VxSAS service for VVR

Table 12-1 Task List: New RDC configuration (Continued)

Objective	Tasks
“Configuring the cluster” on page 355	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring disk groups and volumes” on page 370	<ul style="list-style-type: none"> ■ Planning your storage layout ■ Create disk groups ■ Create volumes ■ Managing disk groups and volumes
“Installing and configuring the application or server role” on page 379	Installing and configuring the application or server role on the cluster nodes
“Configuring the service group” on page 382	<ul style="list-style-type: none"> ■ Using the applicable wizard to create and configure the VCS service group ■ Bringing the service group online
“Creating the primary system zone” on page 433	In the VCS console, selecting the service group and configuring the primary zone nodes as zone 0
“Verifying the installation in the primary zone” on page 434	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“Creating a parallel environment in the secondary zone” on page 435	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Configuring disk groups and volumes for the application, matching the configuration on the primary zone ■ Adding the secondary nodes to the cluster ■ Installing and configuring the application or server role
“Setting up the Replicated Data Sets (RDS)” on page 441	Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
“Configuring a hybrid RVG service group for replication” on page 452	<ul style="list-style-type: none"> ■ Creating a hybrid Replicated Volume Group (RVG) service group ■ Configuring the hybrid RVG service group

Table 12-1 Task List: New RDC configuration (Continued)

Objective	Tasks
“Setting a dependency between the service groups” on page 462	Setting up an online local hard dependency of the application service group (the parent) on the RVG service group (the child)
“Adding the nodes from the secondary zone to the RDC” on page 463	<ul style="list-style-type: none"> ■ Using the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG service group ■ Configuring the secondary nodes as zone 1 ■ Configuring the IP resources for failover
“Verifying the RDC configuration” on page 469	Verifying that failover occurs first within zones and then from the primary to the secondary zone

Reviewing the requirements

Before you install Veritas Storage Foundation HA for Windows, make sure that your configuration meets the following requirements. This replication recovery solution requires installation and configuration at a primary zone and a secondary zone.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 12-2](#) estimates disk space requirements for SFW HA.

Table 12-2 Disk space requirements

Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 5.1 Service Pack 2 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/358407>
- Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:
<http://entsupport.symantec.com/docs/358406>
- Review the operating systems supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported operating systems

SFW and SFW HA have client and server components that run on specific Windows operating systems.

SFW and SFW HA software for servers

Your system must run one of the operating systems listed in the tables to install the SFW or SFW HA server software:

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 12-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required

Table 12-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

The following table shows supported Windows Server 2008 operating systems.

Note: SFW HA supports Windows Server 2008 R2 without Hyper-V x64: Standard Edition, Enterprise Edition, Datacenter Edition.

Table 12-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server Note: SFW HA does not support Hyper-V.	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2

Table 12-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition Note: SFW HA does not support Hyper-V.	
Windows Server 2008 R2 for IA systems	IA64		
Windows Storage Server 2008			

SFW and SFW HA software for clients

The SFW or SFW HA client software can be installed on any one of the operating system versions, editions, and architectures that the Server Components are supported on, except Server Core.

The following table shows the additional client operating systems that are supported.

Table 12-5 Supported client operating systems

Windows Client	Platform	Edition	Service Pack
Windows XP	x86, x64		SP3 required
Windows Vista	x86, x64,	Ultimate Edition, Business Edition, Premium Edition	SP1 or SP2 required
Windows 7	x86, x64	Ultimate Edition, Business Edition, Premium Edition	

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs. See "[Best practices for SFW HA](#)" on page 342.
- NIC teaming is not supported for the VCS private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each application virtual server.
 - A minimum of one static IP address for each physical node in the cluster.

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements for SFW HA

The following permissions are required:

- You must be a domain user.
- You must be a member of the local Administrators group on all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements for SFW HA

Please review the following additional requirements:

- Installation media for all products and third-party applications.
- Licenses for all products and third-party applications.

- You must install the operating system in the same path on all systems. For example, if you install Windows Server on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- For a Replicated Data Cluster, install only in a single domain.

Best practices for SFW HA

Symantec recommends that you perform the following tasks:

- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command. This is applicable for a Replicated Data Cluster configuration.

Note: Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command.

Reviewing the configuration

During the configuration process you will create virtual IP addresses for the following:

- Application virtual server; this IP address should be the same on all nodes at the primary and secondary zones

- Replication IP address for the primary zone
- Replication IP address for the secondary zone

Before you start deploying your environment, you should have these IP addresses available.

Sample configuration

The sample setup has four servers, two for the primary zone and two for the secondary zone. The nodes will form two separate clusters, one at the primary zone and one at the secondary zone.

The procedures in this section are illustrated by a sample deployment and use the following names to describe the objects created and used during the installation and configuration. These names apply to the FileShare application. If you work with a different application, your names will be different.

Table 12-6 RDC configuration objects

Object Name	Description
Primary zone	
SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
FS	File Share server name
FS_SG	File Share service group
FS_SG_DG	Cluster disk group names
FS_REPLOG	Replicator log volume required by VVR
Secondary zone	
SYSTEM3 & SYSTEM4	First and second nodes of the secondary zone
All the other parameters are the same as on the primary zone.	
RDS and VVR Components	
FS_RVG	RVG name for File Share server

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.

Note: Follow the appropriate procedure for verifying the DNS settings, as the procedures for Windows Server 2003 and Windows Server 2008 are slightly different.

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 VCS 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.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 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.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 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 and Dial-up 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 VCS private network.

- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in 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**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 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 VCS 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.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must also select the option to install VVR.

Setting Windows 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.

[Table 12-7](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 12-7 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 12-7 Installation behavior with unsigned drivers (Continued)

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.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see ["Installing Symantec Trusted certificate for unsigned drivers"](#) on page 347.

To change the driver signing options on each local system

- 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, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.

If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

Installing Symantec Trusted certificate for unsigned drivers

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.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

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 HA, 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, select the appropriate SFW product options, and click **Next**.

Client Components	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used while configuring high availability. Required to install the Solutions Configuration Center which provides information and wizards to assist configuration
Volume Replicator (VVR)	If you plan to use Veritas Volume Replicator (VVR) for replication, select this option to install VVR.

- 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.

If you are using multiple paths and selected 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 path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.
---	---

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:	On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).
--	---

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation for Windows (Server Components) with the Microsoft Cluster Server (MSCS) option, 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. For more information, see the *Veritas Storage Foundation Administrator's Guide*.

Review the messages, and then click **OK**.

- 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.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

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 VxSAS

You can run the VVR Security Service Configuration (VxSAS) wizard after you install SFW HA on both the primary and secondary nodes. When you run the wizard, you can then specify the primary and secondary sites in one step.

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server environment, 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 Launch the VVR Security Service Configuration Wizard.

Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.

or

Type `vxsascfg.exe` at the command prompt.

- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name Enter the administrative account name.
(domain\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

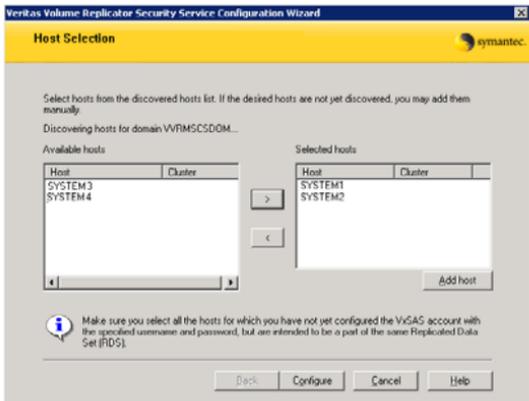
- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. 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 displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with 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.

Configuring the cluster

After you install SFW HA using the installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) does the following:

- Sets up the cluster infrastructure, including LLT and GAB
- Configures Symantec Product Authentication Service in the cluster
- Configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters

Before you create a cluster, do the following:

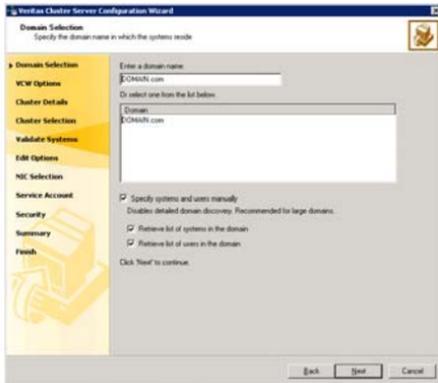
- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

For complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations, see the *Veritas Cluster Server Administrator’s Guide*.

Note: Add only systems in the primary zone (zone 0) to the cluster at this time.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

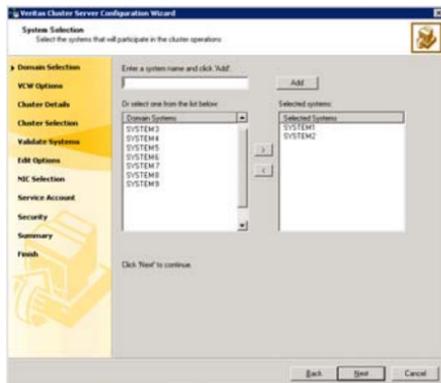


Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 358.
- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.

If you chose to retrieve the list of systems, proceed to [step 6](#) on page 357. Otherwise, proceed to the next step.

- On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster. Proceed to [step 8](#) on page 358.
- On the System Selection panel, specify the systems for the cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

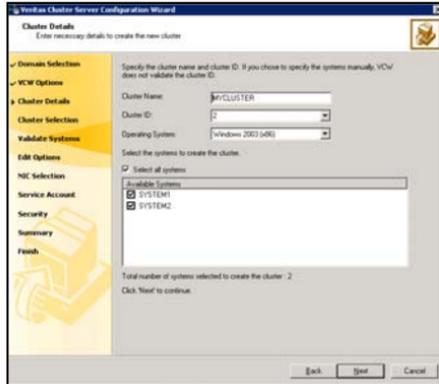
- The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

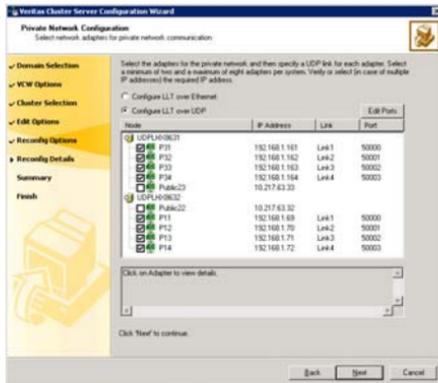
Operating System From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



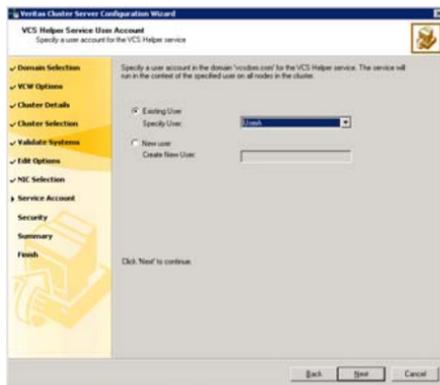
- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

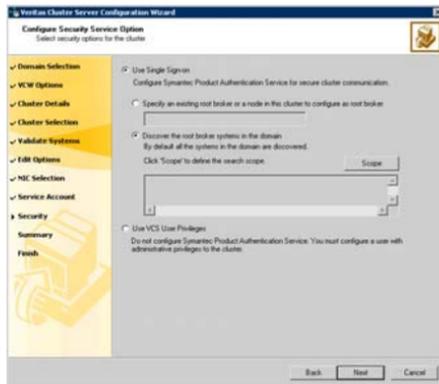
This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:

- Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 356, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:
- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode. For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.
 If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.
 If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.
- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.
 For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 12-8 contains some more examples of search criteria.

Table 12-8 Search criteria examples

1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .

Table 12-8 Search criteria examples

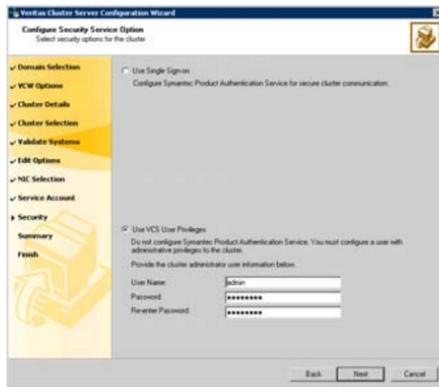
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
 - Click **Next**.
- 14 Review the summary information on the Summary panel, and click **Configure**.
- The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters. To configure the ClusterService group later, click **Finish**.
- At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

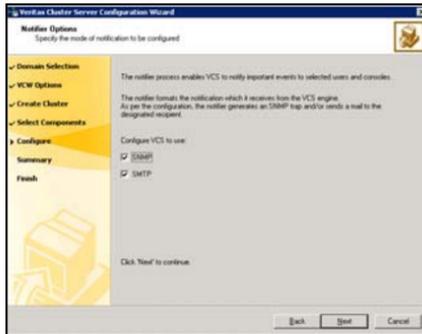
- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See “Configuring notification” on page 366.

Configuring notification

This section describes steps to configure notification.

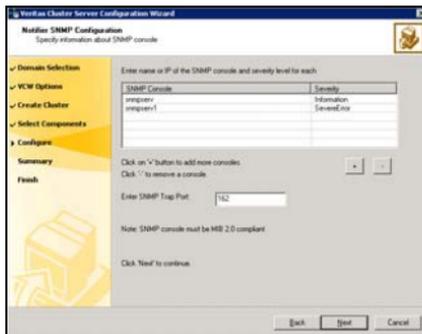
To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.



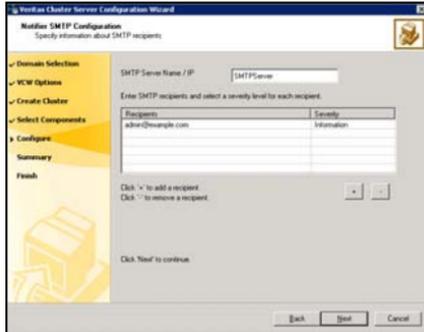
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.



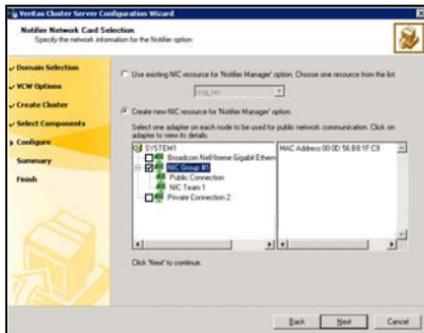
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

Configuring disk groups and volumes

A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts.

Use Veritas Storage Foundation for Windows to create cluster disk groups and dynamic volumes for the application on the shared storage.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - Control Panel - System Settings). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Configuring disk groups and volumes involves the following tasks:

- [“Planning disk groups and volumes”](#) on page 370
- [“Creating dynamic cluster disk groups”](#) on page 372
- [“Creating dynamic volumes”](#) on page 373

Planning disk groups and volumes

The requirements for disk groups and volumes depend on the type of application or server role. Review the requirements and best practices for your application or server role:

- [Planning your File Share storage](#)
- [Planning your IIS storage](#)
- [Planning your Microsoft Virtual Machine storage](#)
- [Planning your storage for additional applications](#)

Planning your File Share storage

Considerations for planning the File Share storage include the following:

- The disk group and volumes for the file server shared directory must be configured on shared storage.
- When you configure a new set up, create the disk groups and volumes on the shared storage first, then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.

Planning your IIS storage

Considerations for planning the IIS storage include the following:

- The disk groups and volumes which will host the directory and files for the web sites must be on the shared storage.
- For a new IIS installation, the directory for the web sites must be created on volumes on the shared storage.
- For existing web sites, stop the sites and then move the web site content to volumes on the shared storage. You must also reconfigure the home directory location for the web site in IIS and then restart the web site again.

Planning your Microsoft Virtual Machine storage

Make sure the volumes that contain the shared virtual disk files for the virtual machines are located on the shared storage.

Planning your storage for additional applications

The information in this section is generic to any application. Make sure that you create the appropriate disk groups and volumes to hold the application data. If your application requires replication of registry keys between the cluster systems, Symantec recommends that you create a dedicated RegRep volume so that its MountV dependency is not linked with any other application-specific resources in the group.

Decide how you want to organize the disk groups and the number and type of volumes you want to create. Considerations include the following:

- The number of disk groups that are needed
The number of disk groups depends on your application and the planned organization of the data. VCS requires that you install the application program files on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application would be contained in a single disk group.
- The type of volumes you want to create.
 - Mirrored and RAID-5 volumes provide fault tolerance for critical data.
 - Striped volumes add performance capability.
 - Volumes that are both mirrored and striped offer both performance and fault tolerance.

Note: If you plan to use replication software, such as VVR, do not use software RAID-5 volumes. This does not apply to hardware RAID-5.

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.
- The implications of backup and restore operations for the disk group setup.
- The sizes of databases and logs, which depend on the traffic load.

Creating dynamic cluster disk groups

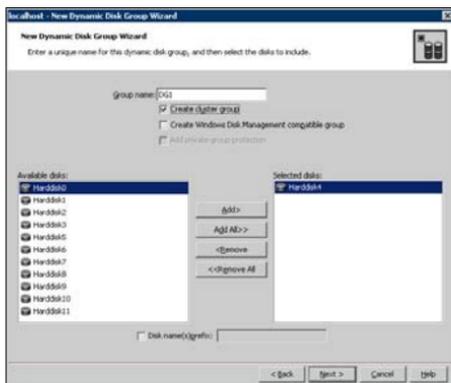
Follow the steps in this section to create one or more disk groups for your application.

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 disk group name (for example, DG1).
- 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.

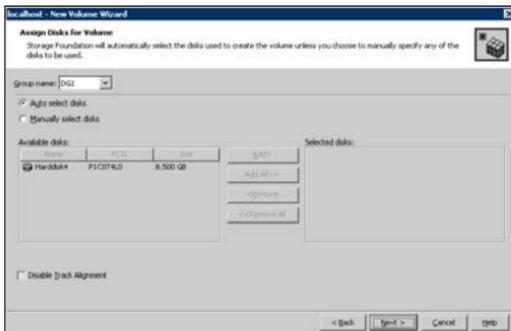
- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

After you create the disk groups, make the disks within them usable by creating the dynamic volumes that will store data.

To create dynamic volumes

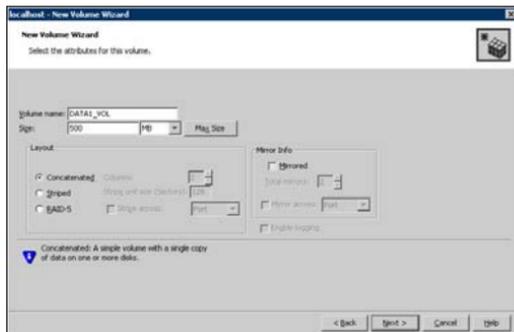
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu 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.
- 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

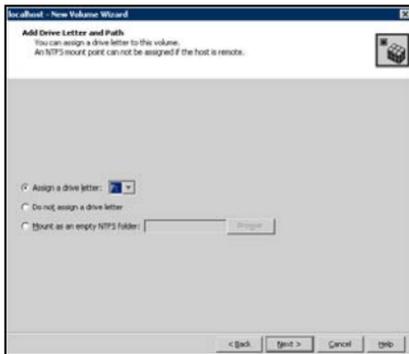
store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the volume attributes.

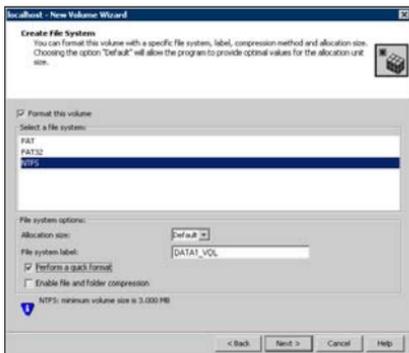


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** 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 additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Managing the disk group and volumes

This section includes the following procedures:

- Importing a disk group and mounting a shared volume
- Unmounting a volume and deporting a disk group

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**.

Installing and configuring the application or server role

This section provides considerations for installing and configuring your application or server role. It includes the following topics:

- [Configuring a File Share server role](#)
- [Configuring a Print Share server role](#)
- [Installing and configuring the IIS application](#)
- [Installing and configuring Microsoft Virtual Server](#)
- [Installing additional applications](#)

Configuring a File Share server role

When you configure a File Share server role, keep in mind the following:

- Configure the disk group and volumes for the file server shared directory on shared storage.
- When you configure a new set up, create the disk group and volumes on the shared storage first, then create the directory structure for the file shares on the shared storage.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage using the practices recommended by Microsoft.
- The FileShare agent is installed automatically with SFW HA.

Configuring a Print Share server role

When you configure a Print Share server role, keep in mind the following:

- Make sure the printer is connected to the network and is configured with an IP address.
- Install software drivers for the network printer on all systems in the cluster.

To add a print driver

- 1 Open the **Printers** Control Panel.
- 2 Click **File > Server Properties**.
- 3 In the Print Server Properties dialog box, click the **Drivers** tab.
- 4 Click **Add**. This launches the Add Printer Driver wizard.

- 5 Follow the wizard instructions to add the printer driver on the system. You must add the driver on each system that will be part of the service group.

Installing and configuring the IIS application

When you install and configure the IIS application, keep in mind the following:

- Install and configure IIS identically on all nodes hosting the service group. The sites to be monitored must be on shared storage.
- Import the cluster disk group and mount the volumes that contain the web site data, on the first node.
- For a new IIS installation, while you are creating new web sites, create the site folder on the shared storage and place the site content in that folder.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. For instructions, see the IIS documentation.
- For existing web sites, stop the sites and then move the web site content to volumes on the shared storage. Reconfigure the home directory location for the web site in IIS and then restart the web site again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group.

To synchronize the IIS configuration on Windows 2003 systems

- 1 Synchronize the IIS configuration on all nodes that will host the IIS service group. Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system.

For example, the following command copies the IIS metabase to *target_system*. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing and configuring Microsoft Virtual Server

When you install and configure the Microsoft Virtual Server, keep in mind the following:

- Install and configure Microsoft Virtual Server identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- If you plan to enable detailed monitoring for the virtual machine resources, install and configure Virtual Machine Additions on each virtual machine.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

Installing additional applications

The following are some very generic points for installing any application:

- Before you install the application, make sure that the disk group and volumes are mounted on the node.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes must be on drive C.
- Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- You should install the data files and any associated files, such as log files, on the shared storage.

Configuring the service group

You can use the Application Configuration Wizard to configure any application for which application specific wizards have not been provided. Depending on the application that you have installed, complete the appropriate procedure to configure the following service groups:

- [Configuring the File Share service group](#)
- [Configuring the Print Share service group](#)
- [Configuring the IIS service group](#)
- [Configuring the MSVirtual Machine service group](#)
- [Configuring the service group for any additional applications](#)

Configuring the File Share service group

To configure the File Share service group, you create a FileShare service group and define the attribute values for its resources. After you create the service group, you must configure the shares to mount automatically at startup.

- Verify that you have local administrator privileges on the system where you run the wizard.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically
`C:\Program Files\Veritas\Cluster Server.`
 - Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
- Verify that the directories to be shared reside on shared drives.
- Mount the drives containing the shared directories on the system where you run the wizard. Unmount the drives from other systems in the cluster.
- Verify that Veritas Command Server service is running on all systems in the cluster.
- Verify that you have the following information ready. The wizard will prompt you for this information:

- A unique virtual computer name to be assigned to the file share server. This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCS`Server` is converted to VCSSERVER.
- A unique virtual IP address to be assigned to the file share server. This is the IP address that the clients can use to access the file share.

Note: Windows Server 2008 does not support accessing file shares using the virtual IP address.

- The list of directories to be shared. The wizard enables you to add existing shares to the VCS configuration. However, you cannot add special shares (shares created by the operating system for administrative and system use). For example, you cannot add the shares ADMIN\$, print\$, IPC\$, and *DriveLetter\$* to the VCS configuration.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

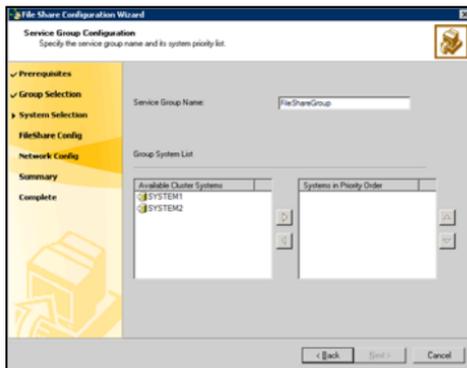
- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).

Creating the File Share service group

For information on resource types, attribute definitions, resource dependencies, and sample service group configurations, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

To configure a FileShare

- 1 Start the File Share Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > File Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the File Share service group.

Available Cluster Systems

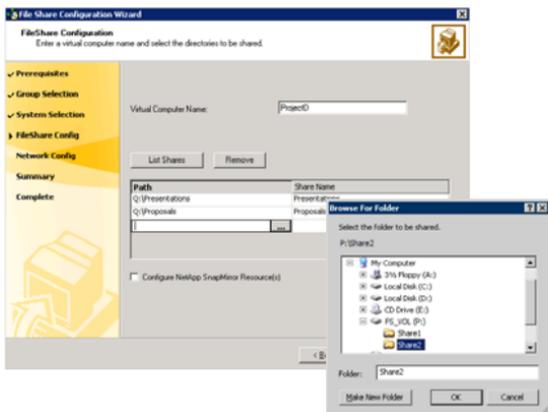
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the File Share Configuration panel, specify the configuration information for the FileShare resources to be created and then click **Next**. The wizard begins validating your configuration. Various messages indicate the validation status.



Virtual Computer Name

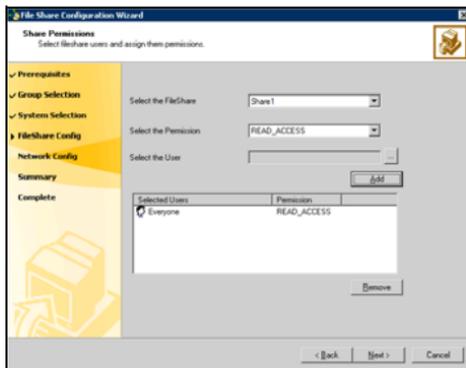
Type a unique virtual computer name by which the server will be known to clients. The virtual name must not exceed 16 characters.

List Shares

Click **List Shares** to view the existing shares on the shared storage, then select a share and click **Add**. You cannot add special shares (shares created by the operating system for administrative and system use).

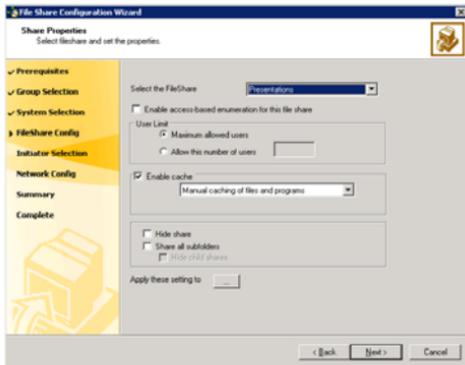
Path	<p>Click the field and either type the path of the directory to be shared or click the ellipsis button (...) to browse for a directory.</p> <p>The selected directories must meet the following conditions:</p> <ul style="list-style-type: none">■ The selected drive, the mount path, and the file path must not exist in the VCS configuration.■ The directories to be shared must reside on shared, non-system drives. <p>The wizard validates the selected directory and displays an error message if the directory does not meet any of the conditions.</p>
Share Name	<p>If a selected directory is already shared, the Share Name column lists the names by which it is shared. You can select a listed share name to make an existing share highly available. You can also create a new share for the same directory by typing a new share name.</p>
Remove	<p>To remove a file share from the configuration, click to select the file share and then click Remove.</p>

- 6 On the Share Permissions panel, select the file share, specify the users for the file shares and assign permissions to them, and then click **Next**.



- Select the FileShare** From the drop-down list, select the file share with which to associate user permissions, or select the default **All FileShares** to set the same permissions for all file shares.
- Select the Permission** From the drop-down list, select the permission to be associated with the user.
- Select the User** Click the ellipsis button (...), select a user, and click **OK**.
- Add** Click **Add** to add the specified user to the Selected Users list. By default, all selected users are given READ_ACCESS permission.
- Selected Users** Displays a list of selected users and the file share permissions. You can configure a maximum of 50 users for each file share. To configure more users, create a user group.
To change the file share permission associated with a user, click a user name in the Selected Users list and then select the desired permission from the Select the Permission drop-down list.
- Remove** To deny file share access to a user, click the user name in the Selected Users list and click **Remove**.

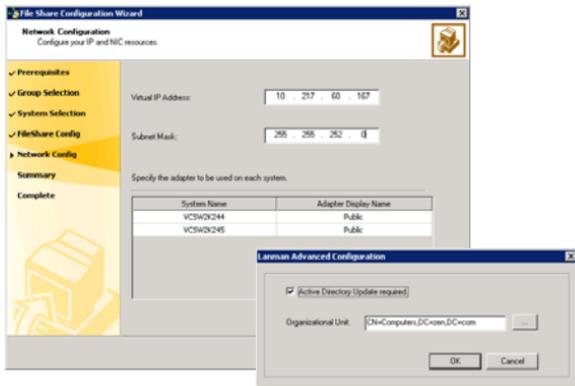
- 7 On the Share Properties panel, set the share properties for the file shares and then click **Next**.



- | | |
|---|---|
| Select the FileShare | From the drop-down list select a file share whose properties you wish to set. |
| Enable access-based enumeration for this file share | Check the Enable access-based enumeration check box to enable the Windows access-based enumeration feature on the selected file share. |
| User Limit | Specify the number of users that are allowed access to the selected file share.
Choose from the following options: <ul style="list-style-type: none">■ Maximum allowed users: Select this option to allow access to the maximum numbers of users allowed on Windows.■ Allow this number of users: Select this option and then type the number of users that you wish to grant access to the selected file share.
If you type zero or a value greater than what Windows supports, access is granted to the maximum allowed users on Windows. |

Enable cache	<p>Check the Enable cache check box to enable local caching of the contents of the selected file share. Then, specify how the contents of the file share are available to users for offline access. In the drop down list select from the following caching options:</p> <ul style="list-style-type: none">■ Manual caching of files and programs: Only the files and programs specified by the user are available offline. This sets the FileShare resource attribute ClientCacheType to MANUAL.■ Automatic caching of programs: All the files and programs that the users access from the file share are available offline. This sets the FileShare resource attribute ClientCacheType to DOCS.■ Optimized automatic caching of files and programs: All files and programs, including executables, are cached locally. The next time the user accesses the executable files, they are launched from the local cache. This sets the FileShare resource attribute ClientCacheType to PROGRAMS.
Hide share	Check the Hide Share check box to make the new share a hidden share.
Share all subfolder	Check the Share all subfolders check box to share the subdirectories.
Hide child shares	Check the Hide child shares check box to hide the shared subdirectories.
Apply these settings to	<p>To apply the specified share properties to multiple file shares simultaneously, do the following:</p> <ol style="list-style-type: none">1 Click the ellipsis (...) button.2 On the Copy Share Properties dialog box, select the file shares from the Available Shares list and click the right arrow to move them to the Selected Shares list. Note that only those files shares that are not already shared are available for selection.3 Click OK. <p>Note: This option is not visible if you are configuring only one share in the service group.</p>

- On the Network Configuration panel, specify information related to your network and then click **Next**.



Virtual IP Address

Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual IP belongs.

Adapter Display Name

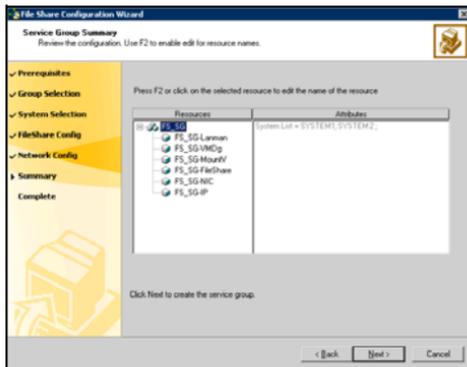
For each system in the cluster, select the public network adapter name.

This field displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow. Verify that you select the adapters assigned to the public network, not the private.

Advanced Settings Click **Advanced Settings** to specify additional details for the Lanman resource.
On the Lanman Advanced Configuration dialog box, complete the following:

- 1 Check **Active Directory Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
 - 2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format
CN=containername,DC=domainname,DC=com.
To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.
By default, the Lanman resource adds the virtual server to the default container "Computers."
 - 3 Click **OK**.
The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- 9 On the Service Group Summary panel, review the service group configuration and click **Next**.
Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- | | |
|------------|---|
| Resources | Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key. |
| Attributes | Displays the attributes and their configured values, for a resource selected in the Resources list. |

- 10 In the completion dialog box, check **Bring the service group online** check box if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the Print Share service group

To configure the Print Share service group, you create a PrintShare service group and define the attribute values for its resources. After you create the service group, you must configure the shares to mount automatically at startup.

- Verify that you have local administrator privileges on the system where you run the wizard.

- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator* mode and then run the VCS commands.

Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Verify that VCS Command Server is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.
- Symantec recommends creating spooler and the replication directories on different disk partitions or volumes.
- Mount the drives with the spooler and the replication directories on the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that the software drivers for the network printers are installed on all systems in the cluster.

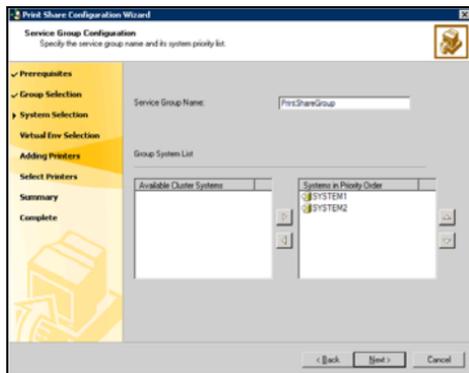
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - A unique virtual computer name to be assigned to the print share server.
This is the name by which clients will access the server. The virtual name must not exceed 16 characters. If you specify a virtual computer name in lowercase letters, the name is converted to uppercase. For example, the name VCSServer is converted to VCSSERVER.
 - A unique virtual IP address to be assigned to the print share server.
 - The network printer's IP address.

Creating the Print Share service group

To create a Print Share service group with a PrintSpool resource

- 1 Start the File Share Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Print Share Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 On the Wizard Options panel, click **Create service group** and click **Next**.

- 4 On the Service Group Configuration panel, specify the service group details and click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the Print Share service group.

Available Cluster Systems

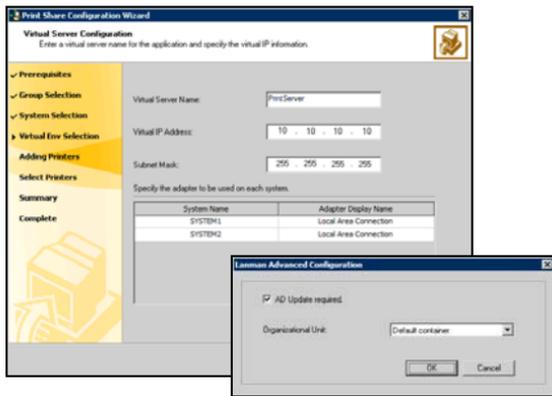
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Virtual Server Configuration panel, specify information related to your network and then click **Next**.



Virtual Server
Name

Type a unique virtual computer name by which the server will be known to clients. Note that the virtual name must not exceed 16 characters.

Virtual IP Address

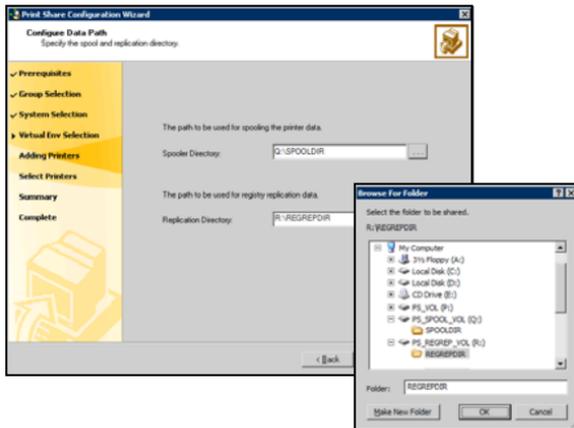
Type a unique virtual IP address for the virtual server.

Subnet Mask

Type the subnet to which the virtual server belongs.

Advanced Settings	<p>Click Advanced Settings to specify additional details for the Lanman resource.</p> <p>On the Lanman Advanced Configuration dialog box, complete the following:</p> <ol style="list-style-type: none">1 Check AD Update required check box to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.2 In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format CN=containername,DC=domainname,DC=com. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."3 Click OK. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
Adapter Display Name	<p>Displays the TCP/IP enabled adapters on a system, including the private network adapters, if applicable. To view the adapters associated with a system, click the Adapter Display Name field and click the arrow.</p> <p>For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.</p>

- 6 On the Configure Data Path panel, specify the spool and registry replication directories and then click **Next**.



Spooler Directory

Type the path or click ... (ellipsis button) to browse for the directory. All print commands will be spooled at this location.

Replication Directory

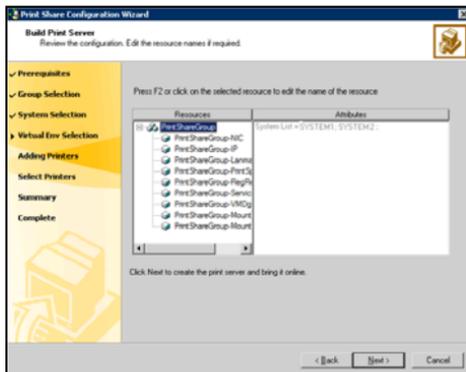
Type the path or click ... (ellipsis button) to browse for the directory. All changes related to the printer registry keys will be logged at this location.

The selected directories must fulfill the following conditions:

- The selected drive, the mount path, and the file path must not exist in the VCS configuration.
- The directories to be shared must reside on shared, non-system drives.
Symantec recommends creating the directories for replication and spooling on different mounts.

- 7 On the Build Print Server panel, review the configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**. The wizard starts running commands to add the PrintSpool resource and the resources on which the

PrintSpool resource depends, including the Lanman and ServiceMonitor resources.



Resources

Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes

Displays the attributes and their configured values, for a resource selected in the Resources list.

8 Bring the PrintSpool resource online.

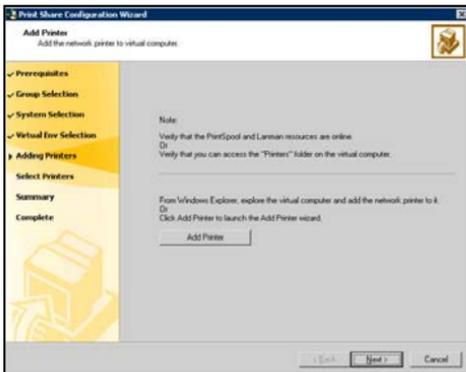
Proceed to the next step to add the network printer to the virtual computer created by the Lanman resource and to create a new TCP/IP port for the printer.

To add the network printer to the virtual computer

- 1 Launch the Add Printer wizard to add the network printer to the virtual computer. Before starting the Add Printer wizard, verify that the PrintSpool and Lanman resources are online in your configuration.

To launch the Add Printer wizard, return to the Print Share Configuration Wizard and click **Add Printer** on the Add Printer panel, or in Windows Explorer, search for the virtual computer, explore the virtual computer by

double-clicking its name and on the virtual computer's Printers folder, double-click **Add Printer**.

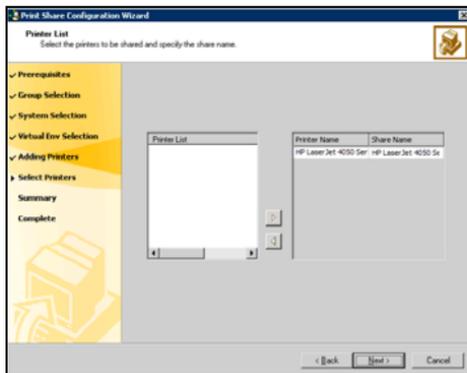


- 2 In the Add Printer wizard, review the information in the Welcome panel and click **Next**.
- 3 Follow the wizard instructions to add the network printer to the virtual computer.
In the Printer Sharing dialog box, always choose the **Do not share this printer** option.
Repeat these steps for each additional printer to be installed.
- 4 Return to the Print Share Configuration Wizard, and proceed to the next step to configure a PrintShare resource in your service group and bring it online.

To configure a PrintShare resource for the service group

- 1 On the Add Printer panel, click **Next**.

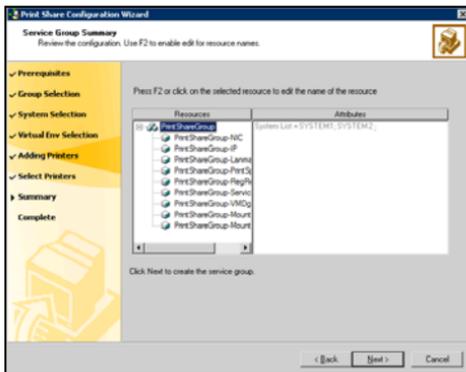
- On the Printer List panel, specify the printers to be included in the Print Share service group and then click **Next**.



- Printer List** Click to select the printer, and then click the right arrow to include the selected printers in your service group. To remove a selected printer from your service group, click the printer from the Printer Name list and click the left arrow.
- Share Name** Type a unique share name for the printer by which it will be known to clients. If you previously chose to share the printer, VCS uses the printer's share name.

- On the Service Group Summary panel, review the service group configuration and then click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration.

Click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



Resources Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

Attributes Displays the attributes and their configured values, for a resource selected in the Resources list.

- 4 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the IIS service group

To configure the IIS service group, you create an IIS service group and define the attribute values for its resources. After you create the service group, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify the sites to be monitored are on shared storage.
- For IIS 7.0 on Windows Server 2008, you must install the following components:
 - IIS 6 Metabase Compatibility
 - IIS 6 WMI CompatibilityThese options are available under Management Tools on the Role Services page of the Add Roles Wizard. These components are required for the IIS agent to function on Windows Server 2008.
For IIS 7.0 on Windows Server 2008 Server Core, you must install IIS in the specified order. See [“Installing IIS 7.0 on Windows Server 2008 Server Core”](#) on page 405 for instructions.
- Do not use the IIS agent to configure SMTP and NNTP sites if you have Microsoft Exchange installed.
- Change the default home directory path for all IIS sites to monitored to a location on the shared storage. See the IIS documentation for instructions.
- Verify the port numbers assigned to IIS sites are not used by other programs.
- Synchronize the IIS configuration on all nodes hosting the service group. See [“Synchronizing IIS configuration on Windows 2003”](#) on page 404 for instructions.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server.`
 - Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- To configure IIS agent on Windows Server 2008 Server Core, you must manually add the required resources and configure the service group. You can perform the manual configuration steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console).

Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Verify that the VCS high availability engine, HAD, is running on the system where you run the wizard.
- Mount the drives containing the shared directories from the system from which you run the wizard. Unmount the drives from other systems in the cluster.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that you have the following information ready. The wizard will prompt you for this information:
 - IIS sites to be monitored.
 - Application pools associated with each site.
 - Port numbers associated with each site.
 - Virtual IP addresses and computer names associated with the sites. The virtual IP addresses and the virtual computer names must have forward and reverse entries in the DNS.
- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Synchronizing IIS configuration on Windows 2003

Complete the following steps.

To synchronize the IIS configuration on Windows 2003 systems

Synchronize the IIS configuration on all nodes that will host the IIS service group.

- 1 Run the script `iiscnfg.vbs`, located at `%systemroot%\System32`. The script copies the IIS metabase from the local system to the target system. For example, the following command copies the IIS metabase to `target_system`. You must enter a valid user name and password for the target system.

```
%systemroot%\System32> iiscnfg /copy /ts target_system /tu  
user_name /tp password
```

- 2 Stop and restart IIS Admin Service on all nodes.

Installing IIS 7.0 on Windows Server 2008 Server Core

Complete the following steps.

To install IIS 7.0 on Windows Server 2008 Server Core

- 1 Type the following at the command prompt:

```
start /w pkgmgr  
/iu:IIS-WebServerRole;IIS-WebServer;IIS-CommonHttpFeatures;  
IIS-StaticContent;IIS-DefaultDocument;IIS-DirectoryBrowsing;  
IIS-HttpErrors;IIS-HttpRedirect;IIS-ApplicationDevelopment;  
IIS-ASP;IIS-CGI;IIS-ISAPIExtensions;IIS-ISAPIFilter;  
IIS-ServerSideIncludes;IIS-HealthAndDiagnostics;  
IIS-HttpLogging;IIS-LoggingLibraries;IIS-RequestMonitor;  
IIS-HttpTracing;IIS-CustomLogging;IIS-ODBCLogging;IIS-Security;  
IIS-BasicAuthentication;IIS-WindowsAuthentication;  
IIS-DigestAuthentication;  
IIS-ClientCertificateMappingAuthentication;  
IIS-IISCertificateMappingAuthentication;  
IIS-URLAuthorization;IIS-RequestFiltering;IIS-IPSecurity;  
IIS-Performance;IIS-HttpCompressionStatic;  
IIS-HttpCompressionDynamic;IIS-WebServerManagementTools;  
IIS-ManagementScriptingTools;IIS-IIS6ManagementCompatibility;  
IIS-Metabase;IIS-WMICompatibility;IIS-LegacyScripts;  
IIS-FTPPublishingService;WAS-WindowsActivationService;  
IIS-FTPPublishingService;IIS-FTPService
```

- 2 Verify that all the components specified in the earlier step have successfully installed. Type the following at the command prompt:

```
notepad C:\windows\logs\cbs\cbd.log
```

This opens the log file, `cbd.log`, in the Notepad text editor.

- 3 Check the entries in the log file, `cbd.log`. The last log entry should resemble the following:

```
Info CBS Pkgmgr: return code: 0x0
```

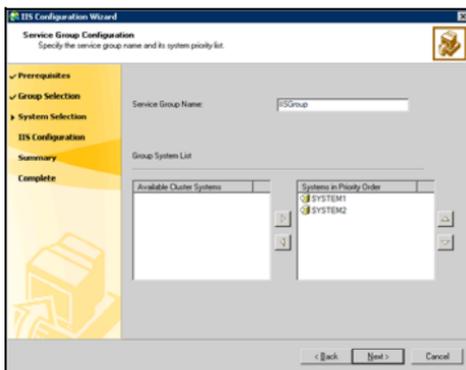
This message indicates that all the components are installed successfully.

- Run the oclist command to verify that the following components are installed:
IIS-WebServerRole; IIS-WebServer; IIS-IIS6ManagementCompatibility;
IIS-Metabase; IIS-WMICompatibility; IIS-FTPPublishingService;
WAS-WindowsActivationService; IIS-FTTPublishingService; IIS-FTPService
Type the following at the command prompt:
oclist
- Repeat the steps on all the nodes on which you wish to configure the IIS service group.

Creating the IIS service group

To create an IIS service group

- Start the File Share Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**.
- Review the information in the Welcome panel and click **Next**.
- On the Wizard Options panel, click **Create service group** and click **Next**.
- On the Service Group Configuration panel, specify the service group details and then click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the IIS service group.

Available Cluster Systems

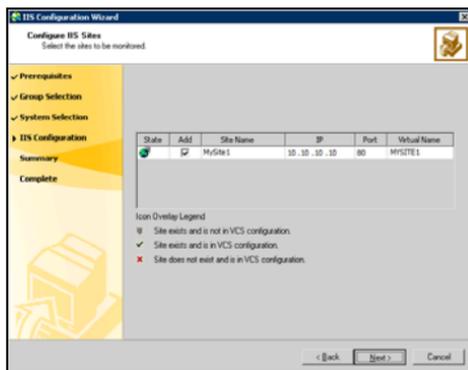
Select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.

To remove a system from the service group's system list, click the system in the Systems in Priority Order box and click the left arrow.

To change a system's priority in the service group's system list, click the system from the Systems in Priority Order and click the up and down arrows.

System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- 5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names, and then click **Next**.

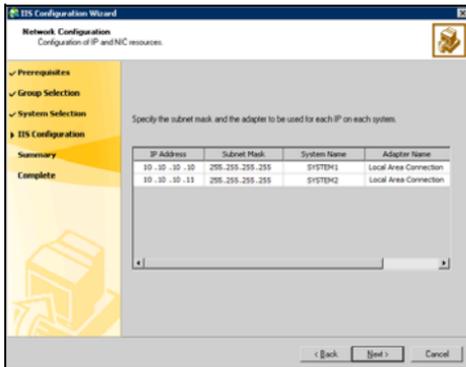


Add Check the check box corresponding to the site to be configured in VCS.

IP Type the virtual IP address for each site to be configured. Make sure that each virtual IP address is associated with only one virtual computer name and vice-versa.

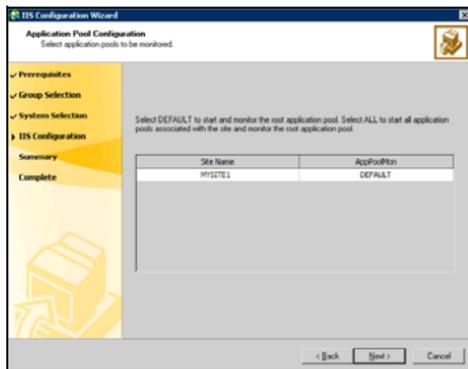
Port Type the port number for each site to be configured.

- 6 On the Network Configuration panel, specify information related to the virtual IP addresses and click **Next**.



- | | |
|--------------|---|
| IP Address | Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems. |
| Subnet Mask | Type the subnet mask associated with each virtual IP address. |
| Adapter Name | Select the adapter associated with the virtual IP address on each system. |

- 7 On the Application Pool Configuration panel, select the monitoring options for application pools associated with each site and click **Next**.



Site Name

Displays the site names.

AppPoolMon

For each site, select the monitoring options from the AppPoolMon list.

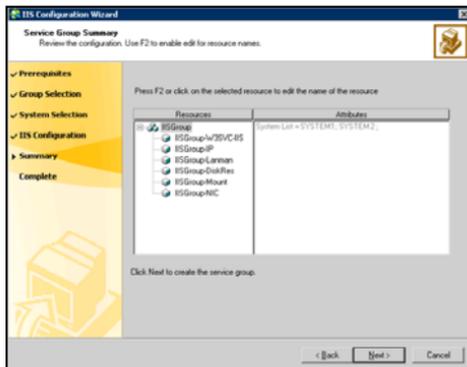
NONE—The agent will not monitor the application pool associated with the site.

DEFAULT—Starts and monitors the root application pool associated with the site.

ALL—Starts all application pools associated with the site and monitors root application pool.

- 8 On the Service Group Summary panel, review the service group configuration and click **Next**. A message appears informing you that the wizard will run commands to modify the service group configuration. Click

Yes. The wizard starts running commands to create the service group. Various messages indicate the status of these commands.



- Resources** Displays a list of configured resources. The wizard assigns unique names to resources. Change the names of resource, if required.
To edit a resource name, select the resource name and either click it or press the F2 key. Edit the resource name and then press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.
- Attributes** Displays the attributes and their configured values, for a resource selected in the Resources list.

- 9 In the completion dialog box, check **Bring the service group online** if you want to bring the service group online on the local system, and then click **Finish**.

Configuring the MSVirtual Machine service group

To configure the MSVirtual Machine service group, you create a MSVirtual Machine service group and define the attribute values for its resources. After you create the service group, you must configure the shares to mount automatically at startup.

Prerequisites

- Verify that you have local administrator privileges on the system where you run the wizard.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service, `%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
- For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the VCS Administrator's Guide for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

- Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.
- Verify that the shared drives required by the applications are mounted.
- Mount the drives containing the shared directories of the virtual machine, on which the wizard will be run. Unmount the drives from other systems in the cluster.

Configuring the service group

- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- Verify that the VCS high availability engine, HAD, is running on the system from which you run the wizard.
- Disable the firewall on each node that will host the service group.
- You must have the following information ready. The wizard will prompt you for this information:
 - The name of the virtual machine.
 - Destination on shared disks for the virtual hard disk files.
 - Network adapters on physical nodes to be associated with network adapters on the virtual machine.
 - Information about monitoring heartbeats (optional).

Creating the MSVirtual Machine service group**To create the MSVirtual Machine service group**

- 1 Start the File Share Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > MSVirtual Machine Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, select the **Create service group** option and click **Next**.
- 4 Enter a name for the service group and specify the systems on which to configure the service group.
 - Enter a name for the service group.
 - In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right arrow to move the systems to the service group's system list.
To remove a system from the service group's system list, click the system in the **Systems in Priority Order** box and click the left arrow.
 - To change a system's priority in the service group's system list, click the system from the **Systems in Priority Order** and click the up and down arrows. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.

- Click **Next**. The wizard then starts validating your configuration. Various messages indicate the validation status.
- 5 Specify details about the virtual machine.
 - Select the virtual machine.
 - For each virtual disk, specify a destination folder where the virtual hard disk files will be moved. Click the Browse icon to browse for folders.
 - To enable detail monitoring for the virtual machine, select the **Monitor Heartbeats** check box and enter failed heartbeat threshold in the **No. of Monitor Cycles** field.

The threshold defines the number of consecutive monitor cycles the agent waits to detect heartbeats from the virtual machine before declaring the resource as faulted.
 - Click **Next**.
 - 6 Select an adapters corresponding to the virtual machine on each system.
 - For each system in the cluster, enter or click a network adapter name to be associated with the network adapters on the virtual machine. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

The fields for the virtual IP address and subnet mask are disabled by design.
 - Click **Next**.
 - 7 Review the service group configuration.

The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.

 - The wizard assigns unique names to resources. Change names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.
 - Click **Next**.
 - A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands.
 - 8 In the completion dialog box, select the check box if you want to bring the service group online on the local system.
 - 9 Click **Finish**.

Configuring the service group for any additional applications

To configure the service group for any additional application, create an application service group and define the attribute values for its resources. You can do this using the Application Configuration Wizard. After you create the service group, you must configure the shares to mount automatically at startup.

Prerequisites

- The application is installed on all the nodes that are going to be part of the service group.
- The shared drives required for the application are mounted on this node.
- The startup type of the application service is set to manual on all nodes in the service group.
- The application service is stopped on all nodes in the service group.
- If your storage is SCSI-3 compliant and you wish to use SCSI-3 persistent reservations, enable SCSI-3 support using Veritas Enterprise Administrator (VEA).
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.

- Port 14141

For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- On Windows Server 2008 Server Core, you have to add the required resources and configure the service group manually. You can perform the steps either directly on the Server Core machine using the command line, or remotely using the Cluster Manager (Java console). Refer to the *Veritas Cluster Server Administrator's Guide* for instructions on using the command line and the Cluster Manager (Java Console).

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

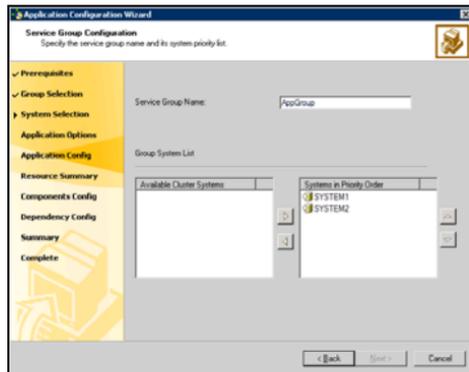
Before configuring the service group, review the resource types and the attribute definitions of the FileShare agent, described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

- Before running the wizard, make sure you have the following information ready:
 - Details (service name, start parameters, startup directory, etc.) of the application that you wish to configure.
 - Shared storage used by the applications.
 - Application registry entries for configuring registry replication.
 - Network and virtual computer (Lanman) details for the application.

Creating the application service group

To create the application service group

- 1 Start the Application Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools >> Application Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.
- 3 In the Wizard Options panel, click **Create service group** and click **Next**.
- 4 Specify the service group name and system list.



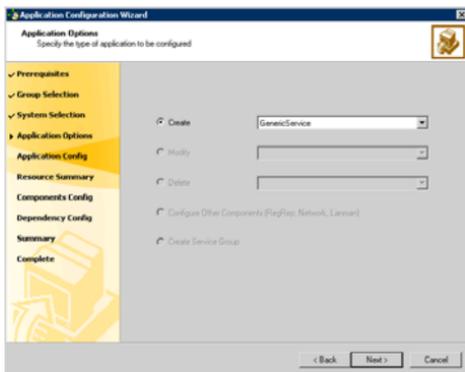
- Enter a name for the service group.

Configuring the service group

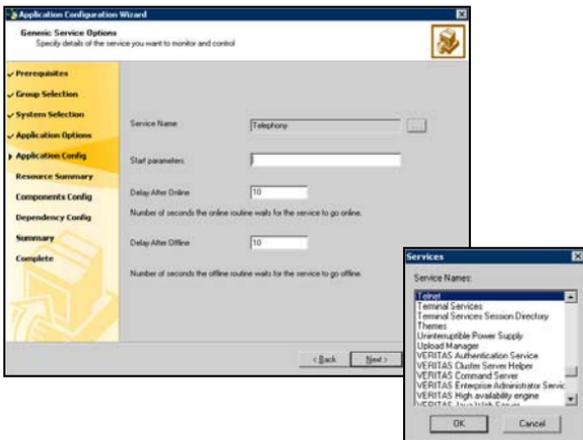
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow icon to move the systems to the service group's system list.
To remove a system from the service group's system list, select the system in the Systems in Priority Order list and click the left arrow.
 - To change a system's priority in the service group's system list, select the system in the Systems in Priority Order list and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 The Application Options panel provides you the option to specify the type of application to be configured. The available options are:
- **Generic Service**: Configures a service using the Generic Service agent. The agent brings services online, takes them offline, and monitors their status. See "[Configuring a GenericService resource](#)" on page 417.
 - **Process**: Configures a process using the Process agent. The agent brings processes online, takes them offline, and monitors their status. See "[Configuring processes](#)" on page 420.
 - **Service Monitor**: Configures a service using the ServiceMonitor agent. The agent monitors a service or starts a user-defined script and interprets the exit code of the script. See "[Configuring a ServiceMonitor resource](#)" on page 424.

Configuring a GenericService resource

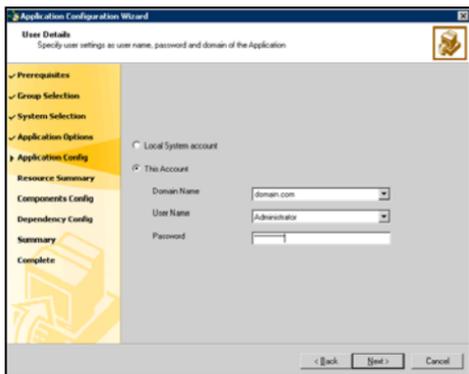
- 1 In the Application Options panel, click **Create**, select **GenericService** from the corresponding drop-down list, and click **Next**.



- 2 Select the service name for which you wish to configure a GenericService resource. Also specify the attributes for the resource.

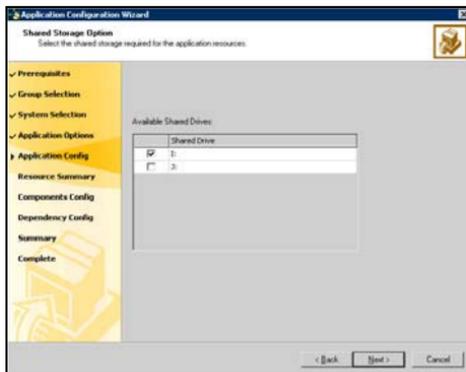


- Click the icon (...) adjacent to the Service Name text box.
 - In the Services dialog box, select a service and click **OK**. The selected service appears in the Service Name text box.
 - In the **Start Parameters** text box, provide the start parameters for the service, if any.
 - In the **Delay After Online** text box, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
 - In the **Delay After Offline** text box, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
 - Click **Next**.
- 3 Specify the information about the user in whose context the service will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the GenericService resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

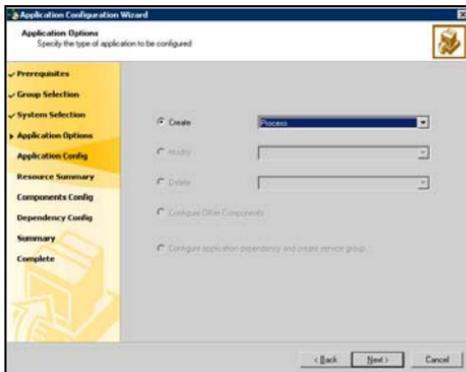


- In the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another GenericService resource, repeat [step 1](#) through [step 5](#).
 - To configure a Process resource, proceed to “[Configuring processes](#)” on page 420 for instructions.
 - To configure a ServiceMonitor resource, proceed to “[Configuring a ServiceMonitor resource](#)” on page 424 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 427 for instructions.

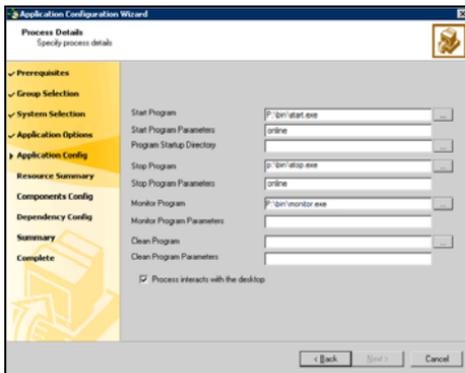
If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 430.

Configuring processes

- 1 In the Application Options panel, click **Create**, select **Process** from the corresponding list, and click **Next**.



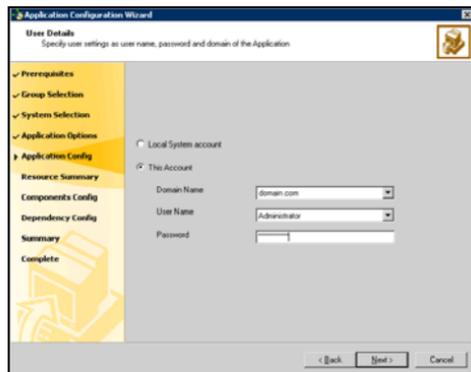
- 2 Specify the details for the process.



- In the **Start Program** text box, specify the complete path of the program that will start the process to be monitored by VCS. You can choose to either type in the location of the program or browse for it using the (...) icon.

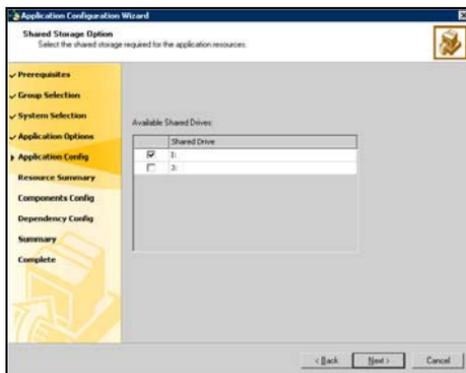
- In the **Start Program Parameters** text box, specify the parameters used by the Process agent start program.
- In the **Program Startup Directory** text box, enter the complete path of the Process agent program or browse for it by clicking the (...) icon.
- In the **Stop Program** text box, enter the complete path of the program that will stop the process started by the Start Program or browse for it by clicking the (...) icon.
- In the **Stop Program Parameters** text box, specify the parameters used by the stop program.
- In the **Monitor Program** text box, enter the complete path of the program that monitors the Start Program or browse for it by clicking the (...) icon.
If you do not specify a value for this attribute, VCS monitors the Start Program. If the Start Program is a script to launch another program, you must specify a monitor program.
- In the **Monitor Program Parameters** text box, specify the parameters used by the monitor program.
- In the **Clean Program** text box, enter the complete path of the Clean process or browse for it by clicking the (...) icon.
- If no value is specified, the agent kills the process indicated by the Start Program.
- In the **Clean Program Parameters** text box, specify the parameters used by the Clean program.
- Select the **Process interacts with the desktop** check box if you want the process to interact with your Windows desktop. Setting this option enables user intervention for the process.
- Click **Next**.

3 Specify information about the user in whose context the process will run.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.
- Click **Next**.

- 4 Select the shared storage required for the Process resource. The shared storage, which you select will be in addition to the mount where the service binaries exist.

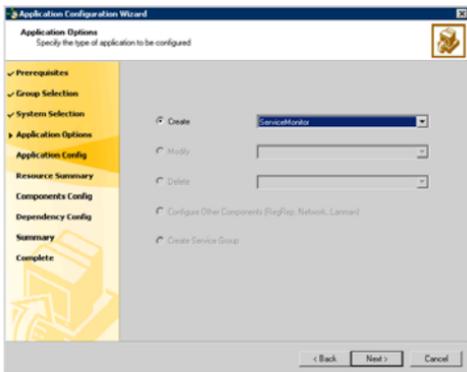


- From the Available Shared Drives box, select the check box adjacent to the shared drive.
 - Click **Next**.
- 5 In the Application Resource Summary panel, review the summary of the Process resource. Click **Back** to make changes. Otherwise, click **Next**.
 - 6 The Application Options panel appears. Select one of the following options:
 - To configure another Process resource, repeat [step 1](#) through [step 5](#).
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 417 for instructions.
 - To configure a ServiceMonitor resource, proceed to [“Configuring a ServiceMonitor resource”](#) on page 424 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to [“Configuring VCS components”](#) on page 427 for instructions.

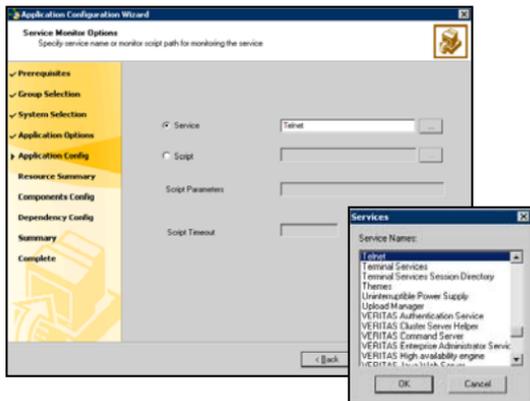
If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 430.

Configuring a ServiceMonitor resource

- 1 In the Application Options panel, click **Create**, select **ServiceMonitor** from the corresponding drop-down list, and click **Next**.



- 2 Specify the service to be monitored or a user-defined script to monitor a service.



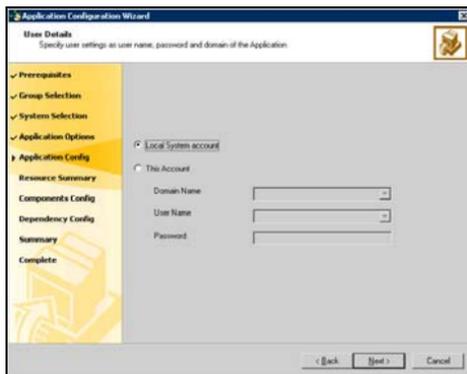
If you want VCS to monitor the service:

- Select the **Service** option and click the icon (...) adjacent to the **Service Name** text box.
- In the Service dialog box, select the service and click **OK**. The selected service name appears in the **Service Name** text box. Alternatively, You may also type in the service name to be monitored.
- Click **Next**.

If you want a script to monitor the service:

- Specify the complete path for the script using the Browse button (...).
- Specify the parameters for the script.
- Specify the time in seconds for the agent to receive a return value from the monitor script.
- Click **Next**.

- 3 Specify the user information in whose context the service will be monitored.



- To configure a service to run in the context of a local system account, click **Local System account**.
- To configure a service to run in the context of another user account, click **This Account**. Specify the **Domain Name**, **User Name**, and **Password** for the user account.

If the service selected in [step 2](#) on page 424 is running in the context of a local system account, the **This Account** option is disabled. Similarly, if the service is running in the context of any other user account, the **Local System account** option is disabled.

- Click **Next**.
ServiceMonitor resource belongs to the category of *persistence* resources. Such resources do not depend on other VCS resources, including shared storage. Hence, the Shared Storage Option panel does not appear if you select the ServiceMonitor option.
- 4 In the Application Resource Summary panel, review the summary of the ServiceMonitor resource. Click **Back** to make changes. Otherwise, click **Next**.
- 5 The Application Options panel appears. Select one of the following options:
 - To configure another ServiceMonitor resource, repeat [step 1](#) through [step 4](#).
 - To configure a GenericService resource, see “[Configuring a GenericService resource](#)” on page 417 for instructions.
 - To configure a Process resource, see “[Configuring processes](#)” on page 420 for instructions.
 - To configure other resources, including FileShare, Registry Replication, and Network resources, proceed to “[Configuring VCS components](#)” on page 427 for instructions.

If you do not want to add any more resources to your service group, proceed to “[Configuring Application Dependencies](#)” on page 430.

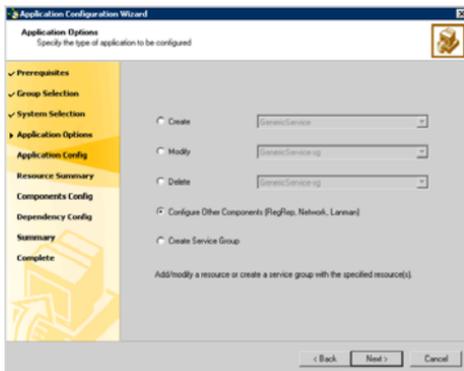
Configuring VCS components

Applications configured using GenericService or Process resources may require network components, or Registry Replication resources. You can configure these VCS components *only* for service groups created using the wizard.

Note: Configure these components only after configuring all application resources. The wizard creates a service group after these components are configured. To add more application resources, you must rerun the wizard in the Modify mode.

To configure VCS components

- 1 In the Application Options panel, click **Configure Other Components**.



- 2 Select the VCS component to be configured for your applications.

The available options are:

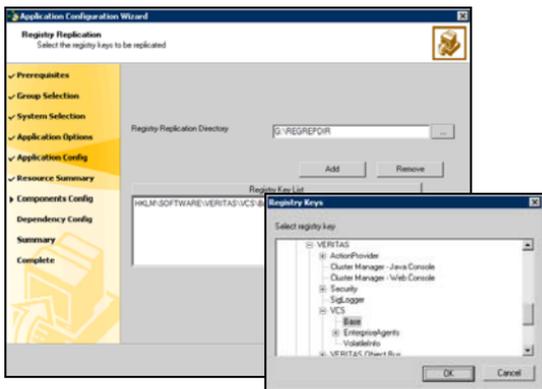
- **Registry Replication Component:** Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to [step 3](#) on page 428.
- **Network Component:** Select this option to configure network components for your application. If you wish to configure a virtual computer name, check **Lanman component** also. To configure a network resource, proceed to [step 4](#) on page 429.

The wizard does not enable the **Lanman Component** check box unless the **Network Component** check box is checked.

To configure Registry Replication

The RegistryReplication panel appears only if you chose to configure the Registry Replication Component in the Application Component panel.

- 3 Specify the registry keys to be replicated.



- Specify the directory on the shared disk in which the registry changes are logged.
- Click **Add**.
- In the Registry Keys dialog box, select the registry key to be replicated.
- Click **OK**.
- The selected registry key is added to Registry KeyList box. Click **Next**.

If you chose Network Component from the Application Component panel, proceed to the next step. Otherwise, proceed to [step 5](#) on page 429.

To configure network components

The Virtual Computer Configuration panel appears only if you chose to configure the Network Component in the Application Component panel.

- 4 Specify information related to your network.
 - Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 16 characters.

Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component panel.
 - Enter a unique virtual IP address for the virtual server.
 - Enter the subnet to which the virtual server belongs.
 - Click **Advanced** to specify additional details for the Lanman resource.
 - Check **AD Update required** to enable the Lanman resource to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
 - In the Organizational Unit field type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box for the virtual server. By default, the Lanman resource adds the virtual server to the default container "Computers."

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **OK**.
 - For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Display Name** field and click the arrow.

Note that the wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private.
 - Click **Next**.
- 5 The Application Options panel is displayed. Select one of the following options:

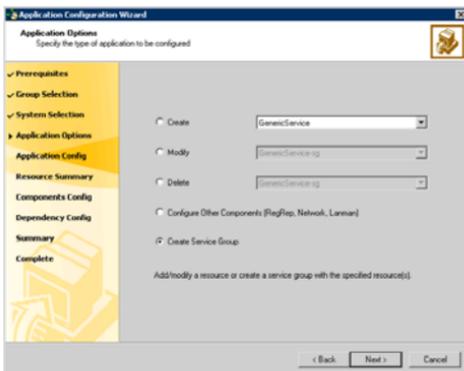
- To configure additional VCS components, repeat [step 1](#) on page 427 through [step 4](#) on page 429.
 - To configure a GenericService resource, see [“Configuring a GenericService resource”](#) on page 417 for instructions.
 - To configure a Process resource, see [“Configuring processes”](#) on page 420 for instructions.
 - To configure a Service Monitor resource, see [“Configuring a ServiceMonitor resource”](#) on page 424 for instructions.
- If you do not want to add any more resources to your service group, proceed to [“Configuring Application Dependencies”](#) on page 430.

Configuring Application Dependencies

The Application Configuration Wizard enables you to create service group for the application resources and other VCS components configured using the wizard. This section describes how to create the service group using the wizard.

To create a service group

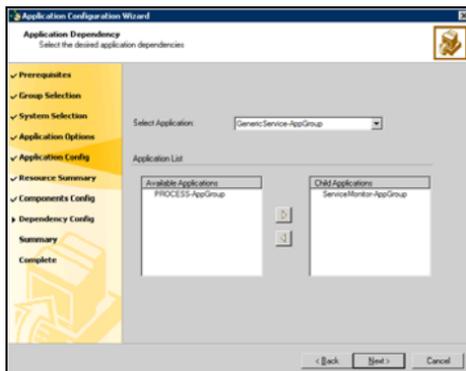
- 1 In the Application Options panel, click **Configure application dependency and create service group**.



The option is enabled only if:

- resources and VCS components are already configured using the wizard.
- you clicked **Modify Service Groups** in the Wizard Options panel.

2 Specify the dependency between the applications.

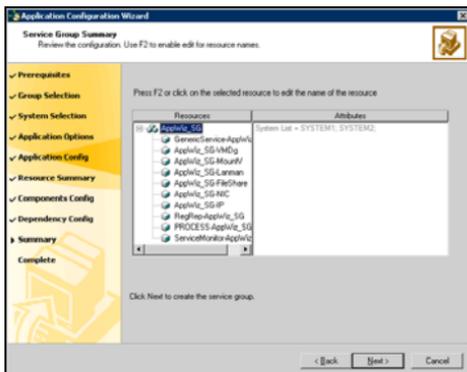


You must have at least two resources configured for Application Dependency panel to appear. Of the two resources, one should either be a GenericService or a Process resource.

- From the Select Application list, select the application that would depend on other applications. The selected application becomes the parent application.
- From the Available Applications list, select the application on which the parent application would depend and click the right-arrow icon to move the application to the Child Applications list.
- To remove an application from the Child Applications list, select the application in the list and click the left arrow.
- Repeat these steps for all such applications for which you want to create a dependency.
- Click **Next**.

The Application Dependency panel enables you to link resources configured using the wizard. If these resources are dependent on other services outside the VCS environment, you should first configure resources for such services and then create the appropriate dependency.

3 Review the service group configuration.



The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.

- The wizard assigns unique names to resources. Change names of resource, if required.

To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press the Esc key.

- Click **Next**.
- A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes**.

The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.

- 4 In the completion panel, check **Bring the service group online** if you want to bring the service group online on the local system.
- 5 Click **Finish** to create the service group and exit the Application Configuration Wizard.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the left pane and the Properties tab in the right pane, select the service group.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.



- 7 Click **OK**.

Verifying the installation in the primary zone

An important part of configuration testing is simulating a failover.

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.
- Simulate a local cluster failover by shutting down an active cluster node.

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in [step 1](#).
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.
The service group you selected is taken offline and brought online on the node that you selected.

Creating a parallel environment in the secondary zone

Before you begin to configure the secondary zone, offline the following resources in the application service group in the primary zone:

- Application resource
- Application virtual IP resource

The remaining resources should be online, including the VMDg resources and the MountV resources.

In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

After you set up a SFW HA environment in the primary zone (zone 0), use the guidelines in the following list to complete the same tasks in the secondary zone (zone 1).

- [“Reviewing the requirements”](#) on page 336
- [“Reviewing the configuration”](#) on page 342
- [“Configuring the storage hardware and network”](#) on page 343
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 346
- [“Configuring VxSAS”](#) on page 352
- [“Adding the systems in the secondary zone to the cluster”](#) on page 436
- [“Configuring disk groups and volumes”](#) on page 370

During the creation of disk group and volumes for the secondary zone, make sure the following are exactly the same as the cluster at the primary zone:

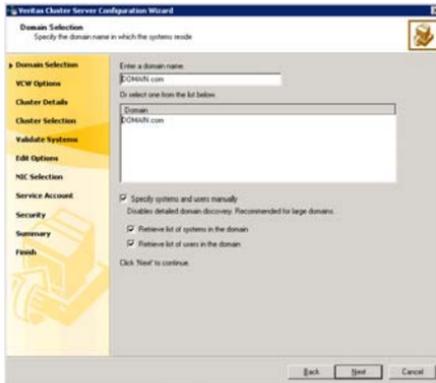
- Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing and configuring the application or server role”](#) on page 379

Adding the systems in the secondary zone to the cluster

Use the following procedure to add the nodes in the secondary zone to the existing cluster.

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



Do one of the following:

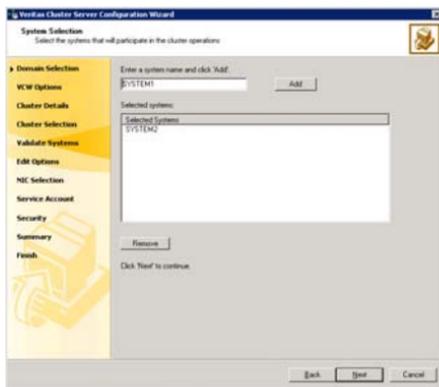
- To discover information about all the systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 439.

- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.

If you chose to retrieve the list of systems, proceed to [step 6](#) on page 438. Otherwise proceed to the next step.

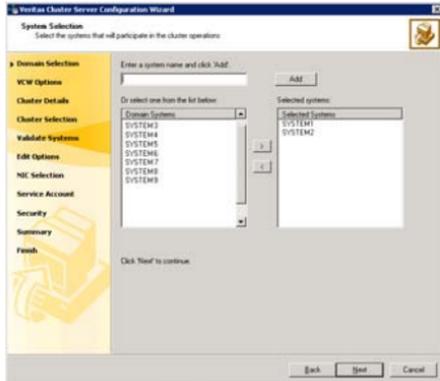
- 5 On the System Selection panel, complete the following and click **Next**.



- Type the name of a node in the cluster and click **Add**.
- Type the name of the system to be added to the cluster and click **Add**. If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

Proceed to [step 8](#) on page 439.

- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.



Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

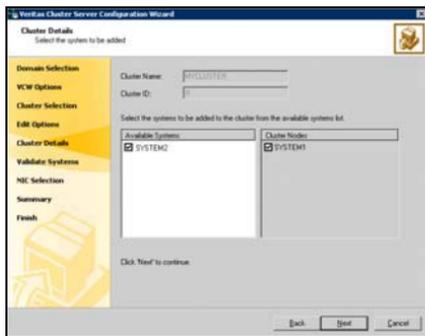
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#), only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**. In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.



The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on how it is configured in the cluster. If LLT is configured over ethernet, you

have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have to use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard will configure the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select **Configure LLT over UDP**.
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links.
Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.

The IP address is used for the VCS private communication over the specified UDP port.

- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard will configure the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 14 On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
- 17 The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

Return to the task list “[Creating a parallel environment in the secondary zone](#)” on page 435.

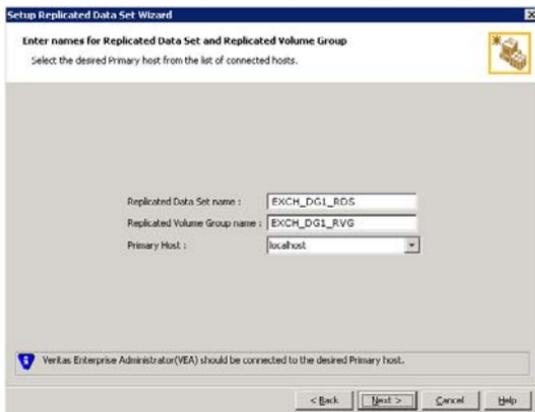
Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- Verify that the data volumes are not of the following types. VVR does not support the following 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
 - Volumes names containing a comma
- Verify that the cluster disk group is imported and the volumes are mounted in the primary and secondary zone.
- Configure the VxSAS service if you have not already done so. See “[Configuring VxSAS](#)” on page 352.

To create the Replicated Data Set

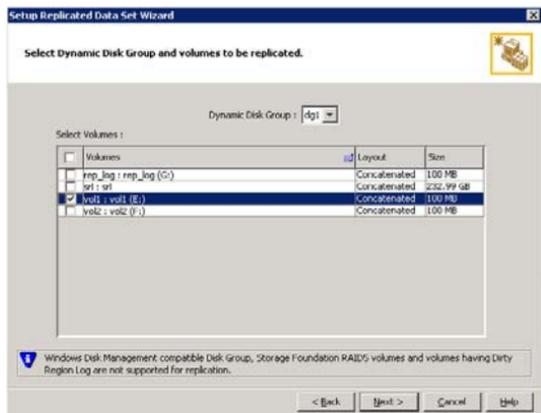
- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the 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 (APP_REPL_LOG). 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. <ul style="list-style-type: none">■ 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:

The screenshot shows a window titled "Setup Replicated Data Set Wizard" with a sub-header "Edit replication settings". Below the sub-header is the instruction "Edit replication settings or click next." and a small icon of a server rack. The main area contains the following fields:

Primary side IP	10.217.53.214
Secondary side IP	10.217.53.215
Replication Mode	Synchronous Override
Replicator Log Protection	AutoDCM
Primary RLINK Name	Pri_RLINK
Secondary RLINK Name	Sec_RLINK

Below these fields is an "Advanced" button. At the bottom left, there is a blue information icon and the text "DHCP addresses are not supported by VWR." At the bottom right, there are four buttons: "< Back", "Next >", "Cancel", and "Help".

- 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	Select the required mode of replication: Synchronous Override , Synchronous , or Asynchronous . The default is synchronous override. 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. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. 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. 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 MISSING.
Replicator Log Protection	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.

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.

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.

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/EAT32 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.

If you have set up additional disk groups for the application, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 441 for each additional disk group. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name.

Configuring a hybrid RVG service group for replication

The RVG service group is a hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For more information about service group types, see the *Veritas Cluster Server Administrator's Guide*.

You configure the RVG service group resources manually by copying and modifying components of the application service group. You then create new RVG resources and bring them online.

Table 12-9 lists the resources in the RVG service group for RDC.

Table 12-9 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the disk group	Volume Manager disk group for the application
VvrRvg for the disk group	Replicated volume group for the application

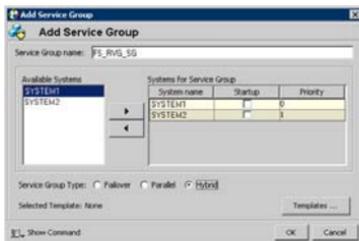
Creating the RVG service group

To contain the resources for replication, you need to create a hybrid replicated volume (RVG) service group.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.

- 3 In the **Add Service Group** window, do the following, in the order presented:



- Enter a name for the service group. Choose a service group name that is meaningful for the type of application you are using. This example adds a FileShare RVG service group called FS_RVG_SG.
- Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.
- Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group's resources manually for RVG by completing the tasks in [Table 12-10](#):

Table 12-10 RVG service group configuration tasks

Task	For more information, see
Copy IP and NIC resources of the application service group, paste and modify them for the RVG service group.	“Configuring the IP and NIC resources” on page 454
Copy the VMDg resources for the disk groups in the application service group, paste and modify them for the RVG service group.	“Configuring the VMDg resources for the disk groups” on page 455
Create the VVR RVG resources for the disk group and enter the attributes for the disk group and the replication IP address.	“Adding the VVR RVG resources for the disk groups” on page 457

Table 12-10 RVG service group configuration tasks

Task	For more information, see
<p>Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk group.</p> <p>Configure the RVG service group's VMDg resources to point to the disk group that contain the RVGs.</p>	<p>"Linking the VVR RVG resources to establish dependencies" on page 458</p>
<p>Delete the VMDg resources from the application service group, because they depend on the replication and were configured in the RVG service group.</p>	<p>"Deleting the VMDg resource from the application service group" on page 458</p>

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC:

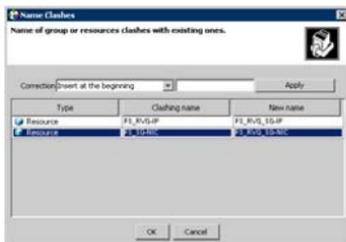
Table 12-11 IP and NIC resources

Resource	Attributes to modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the application service group in the left pane.
- 2 On the **Resources** tab, right-click the IP resource, and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group.
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.

- In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group and click **OK**.



To modify the IP resource and NIC

- In the **Resources** tab display area, right-click the IP resource and select **View > Properties View**.
- In the **Properties View** window, for the **Address** attribute, click **Edit**.
- In the **Edit Attribute** window, enter the VVR IP address for the Primary Zone as the scalar value.
- Close the **Properties View** window.

To enable the IP resource and NIC

- In the **Resources** tab display area, right-click the IP resource and select **Enabled**.
- In the **Resources** tab display area, right-click the NIC resource and select **Enabled**.

Configuring the VMDg resources for the disk groups

Create the VMDg resources in the application service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the application service group for the MountV resource:

Table 12-12 MountV resources

Resource	Attributes to modify
Resources for the disk group for the application files:	

Table 12-12 MountV resources

Resource	Attributes to modify
MountV	VMDg Resource Name Volume Name

To create a VMDg resource for the disk group

- 1 In the VCS Cluster Explorer window, select the application service group in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the disk group, and click **Copy > Self**.
- 3 In the left pane, select the RVG service group.
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the application service group

- 1 In the VCS Cluster Explorer window, select the application service group in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the application and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the volume created for the application.
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created.
- 6 Close the **Properties View** window.

- 7 Repeat the steps for any additional MountV resources for the application.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group.
- 2 In the **Resources** tab display area, right-click the VMDg resource and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add a VvrRvg resource for replication of the disk group. If the application has multiple disk groups, create a separate VvrRvg resource for each disk group. Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 12-13 VvrRvg resources

Resource	Attributes to modify
Resources for the disk group for the application files:	
VvrRvg	VMDgResName IPResName

To create the VVR RVG resource for the disk group

- 1 In the left pane, select the RVG service group. Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window, do the following, in the order presented:
 - Enter the **Resource Name** for the VVR RVG resource.
 - Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the RVG that is being managed.
- 5 Click **OK**.
- 6 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.
- 7 In the **Edit Attribute** window, enter the name of the disk group containing the RVG.
- 8 Click **OK**.

- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name of the IP resource managing the IP address for replication.
- 11 Click **OK**.
- 12 In the **Add Resource** window, verify that the attributes have been modified.
- 13 Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources. Depending on the application you use, your resource names may be different.

Table 12-14 Dependencies for VVR RVG resources for RDC

Parent	Child
Resources for the disk group for the application system files:	
FS_RVG_SG-VvrRvg	The IP for replication, for example FS_RVG_SG-IP
FS_RVG_SG-VvrRvg	The VMDg for the application, for example FS_RVG_SG-VMDg

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group.
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example, FS_RVG_SG_VvrRvg.
- 4 Click the child resource, for example, FS_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources.
Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg resource from the application service group

The VMDg resources must now be manually deleted from the application service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the application group

- 1 In the VCS Cluster Explorer window, select the application service group from the left pane.
- 2 In the **Resources** tab display area, right-click the VMDg resource for the disk group and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the **Resources** tab display area, right-click the VMDg resource for any additional group and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

For each application disk group, add a resource of type RVGPrimary to the application service group and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the application service group for the RVG Primary resources:

Table 12-15 RVG Primary resources

Resource	Attributes to Modify
Resources for the disk group for the application:	
RVGPrimary	RvgResourceName

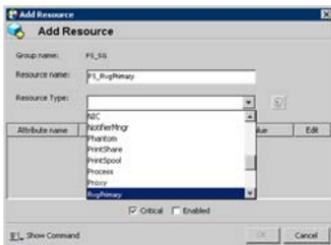
Creating the RVG Primary resources

Create an RVG Primary Resource for replication.

To create the RVG Primary resource for an application's disk group

- 1 In the VCS Cluster Explorer window, right-click the application service group in the left pane, and select **Add Resource**.

- 2 In the **Add Resource** window, do the following, in the order presented:



- Enter the **Resource Name** for the RVG Primary resource for the application disk group.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
 - 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, and click **OK**.
 - 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server 5.1 Administrator's Guide* for more information about the RVG Primary agent.
 - 6 Verify that **Critical** and **Enabled** are both checked.
 - 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the application service group to establish the dependencies between the resources for replication.

Start from the top parent and link the resources for your application.

As an example, [Table 12-16](#) lists the Parent and Child relationship for the FileShare application.

Table 12-16 Dependencies for the RVG Primary resources for RDC

Parent	Child
FS_SG-MountV	FS_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the application service group.
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource.
- 4 Click the child resource.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link any additional RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the application service group online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the application service group.
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource, and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right click any additional RVG Primary resource, and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG service group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.

Setting a dependency between the service groups

The RVG service group must be online on both the primary and secondary zones. However, if a failover occurs from one node to another within the same zone, the RVG service group must fail over along with the application service group.

To ensure that the application service group and the RVG service group fail over and switch together from one node to another within the same zone, set up an online local hard dependency between the RVG service group and the application service group. The application service group depends on the RVG service group.

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the **Service Groups** tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the application service group (the parent service group).
- 5 Click the RVG service group (the child resource).
- 6 In the **Link Service Groups** window, do the following, in this order:
 - Select the **Relationship** of **online local**.
 - Select the **Dependency Type** of **hard**.
 - Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box, do the following, in the order presented:
 - Click **Modify an existing replication service group**. The existing replication service group is selected, by default.
 - Click **Next**.
- 4 If a VCS notice message appears, asking if you want to continue, click **Yes**.
- 5 Specify the system priority list. Do the following in the order presented:
 - In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 - To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
 - Click **Next**.
- 6 If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
- 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.

- 8 In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
- 9 If a VCS error appears, click **OK**.
- 10 In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.
- 11 Review the summary of the service group configuration. In the **Resources** box, click a resource to view its attributes and their configured values in the **Attributes** box.
- 12 Click **Next** to modify the replication service group.
- 13 At the prompt, click **Yes** to modify the service group.
- 14 Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

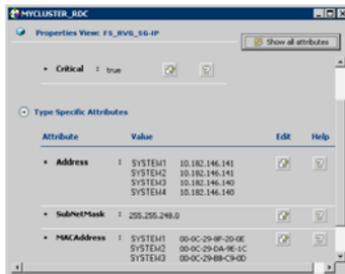
Configuring the IP resources for failover

Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.

If the application fails, VCS tries to fail over the application service group to another system within the same RDC system zone. However, if VCS cannot find a failover target node in the primary zone, it switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the Address attribute. Do the following in the order presented:
 - Select **Per System**.
 - Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - Repeat for all nodes in the primary zone.
 - Select the first node in the secondary zone and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone have the same IP address, that is different from the IP address for the primary zone.



- 6 Close the Properties View window.

Adding the nodes from the secondary zone to the RDC

- 7 Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding the nodes from the secondary zone to the application service group

Use the wizard appropriate to your application to add the nodes from the secondary zone to the application service group.

To add nodes from the secondary zone to the application service group

- 1 Select **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > File Share Configuration Wizard**.
- 2 The File Share Configuration Wizard screen appears. This screen displays the pre-requisites and User Input Required. Verify that you have meet the prerequisites listed and click **Next**.
- 3 The Wizard Options window is then displayed. In the Wizard Options window, select the Modify service group option. Next, select the File Share service group below this option and click **Next**.
- 4 The Service Group Configuration window then displays. In the Service Group Configuration window, review and if required update the system priority list. On the Service Group Configuration window, select the nodes in the secondary zone, use the arrow button to move them from **Available Cluster Systems** to **System in Priority Order**. To change the priority of a system in the **Systems in Priority Order** list, select the system and click the up and down arrow icons. Arrange the systems in priority order in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online. This set of nodes selected for the application service group must be the same as the nodes selected for the RVG service group. Ensure that the nodes are also in the same priority order.
- 5 Click **Next**.
- 6 In the FileShare Configuration window, enter a virtual computer name and select the directories to be shared. Click **Next**.
- 7 In the Share Permissions window, select the fileshare users and assign them permissions. Click **Next**.
- 8 In the Network Configuration window, configure your Virtual IP address and Subnet Mask. Specify the adapter to be used on each system, by entering the system name and adapter display name. Click **Next**.
- 9 In the Service Group Summary window, review the service group configuration. Use F2 to enable edit for resource names. After your review, click **Next**.

- 10 A message appears if the configuration is currently in the Read Only mode. Click **Yes** to make the configuration read and write enabled. The wizard validates the configuration and modifies it.
- 11 Click **Finish**.

Configuring the zones in the application service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the application service group

- 1 From VCS Cluster Explorer, in the left pane, select the application service group.
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- [Bringing the service group online](#)
- [Switching online nodes](#)

Bringing the service group online

After completing all configuration, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the Exchange Server service group online in the primary zone.

To bring the Exchange service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange Server service group (EVS1_SG1).
- 2 Click **Online**.

Switching online nodes

An important part of configuration testing is simulating a failover. Test the failover by switching the application service group between online nodes.

The RVG service group is online in both the primary and secondary zone. However, within a zone, if more than node is configured, the RVG service group should fail over with the application service group.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

Switch the application service group between nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to another.

If you enter the system name manually from the Java Console, specify the name in upper case.

To switch online nodes

- 1 Open the Veritas Cluster Manager (Java Console). Click **Start > All Programs > Symantec > Veritas Cluster Manager (Java Console)**.

Additional instructions for GCO disaster recovery

- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 In the left pane, right-click the service group, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm that you want to switch the service group to the other node.

Additional instructions for GCO disaster recovery

After you complete the tasks for setting up a replicated data cluster for an application service group, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, if a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment. The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up VVR replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See [“Deploying disaster recovery: New application installation”](#) on page 481.

Disaster Recovery

This section includes the following chapters:

- [Disaster recovery: Overview](#)
- [Deploying disaster recovery: New application installation](#)
- [Testing fault readiness by running a fire drill](#)

Disaster recovery: Overview

This chapter covers the following topics:

- [“About a disaster recovery solution”](#) on page 474
- [“Need for implementing a disaster recovery solution”](#) on page 475
- [“Overview of the recovery process”](#) on page 476
- [“Components of VVR that enable disaster recovery”](#) on page 477

About a disaster recovery solution

A disaster recovery (DR) solution is a series of procedures used to safely and efficiently restore application data and services in the event of a catastrophic failure. A typical disaster recovery solution requires that you have a source host on the *primary* site and a destination host on the *secondary* site. The application data is stored on the primary site and replicated to the secondary site by using a tool such as the Veritas Volume Replicator. The primary site provides data and services during normal operation. If a disaster occurs on the primary site and its data is destroyed, a secondary host can take over the role of the primary host to make the data accessible. The application can be restarted on that host.

This chapter is an overview of the Veritas Volume Replicator disaster recovery solution that can be used with SFW HA and VCS. SFW HA also supports array-based hardware replication. SFW HA provides a configuration wizard for disaster recovery, which can be used with either VVR or hardware replication.

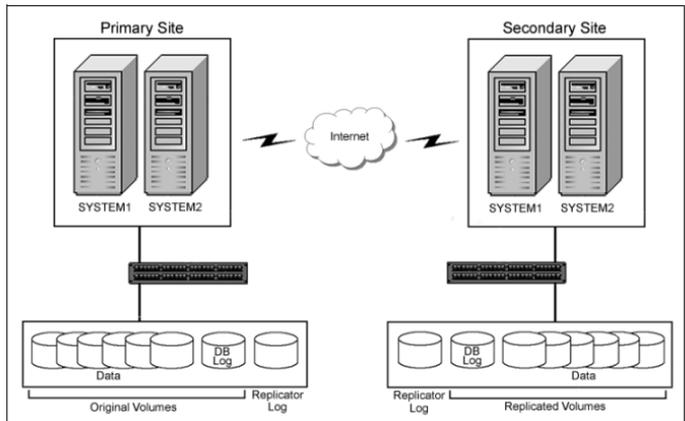
For details on configuring SFW HA disaster recovery using the wizard, see [“Deploying disaster recovery: New application installation”](#) on page 481.

For a SFW and VVR configuration with Microsoft clustering, see the following:

- [Chapter 21, “Deploying SFW and VVR with MSCS” on page 759.](#)
This solution is deployed on Windows Server 2003.
- [Chapter 22, “Deploying SFW and VVR with Microsoft failover clustering” on page 805](#)
This solution is deployed on Windows Server 2008.

The illustration below shows the SFW HA-VVR configuration with VCS. The example has one disk group on each site for the application. Note that a VVR Replicator Log is needed on each site. If there is more than one disk group, an additional Replicator Log is required for each disk group.

Figure 13-1 SFW HA-VVR configuration with VCS



Need for implementing a disaster recovery solution

Two major trends affecting businesses today are reliance on data and geographic distribution. Continuous, consistent, fast, and reliable access to data is important. If a disaster occurs, quick availability of data becomes important. One of the ways of achieving this is by using replication.

A well-designed disaster recovery solution prepares a business for unexpected disasters and provides the following benefits in the event of a disaster:

- Minimizes economic loss due to the unavailability or loss of data.
- Ensures safe and efficient recovery of data and services.
- Minimizes decision making during the disaster recovery.
- Reduces reliance on key individuals.
- Minimizes data loss during recovery and ensures availability of the most recent data.

A strategic disaster recovery (DR) solution can provide businesses with ways to meet their service level agreements, comply with government regulations, and minimize their business risk.

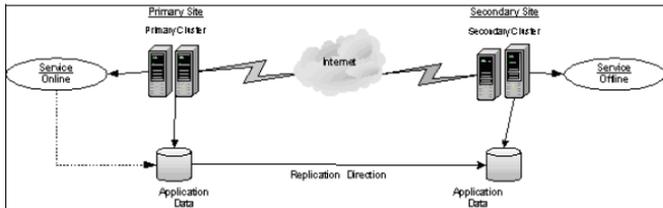
Overview of the recovery process

The illustrations that follow show the typical disaster recovery setup before and after a disaster.

In the illustration before the disaster, the primary host replicates its application data to the secondary host. In a disaster recovery 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.

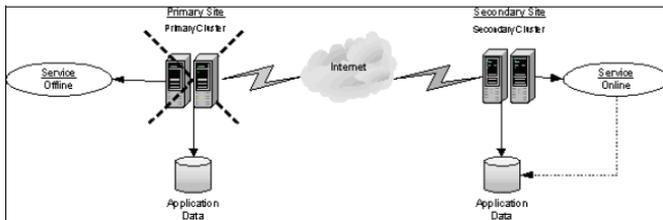
Note that the primary and the secondary sites have clusters to make both the application and VVR highly available.

Figure 13-2 Typical disaster recovery configuration setup



If a disaster, such as an earthquake, causes a failure at the primary site, a host on the secondary site can take over the role of the primary host to make the data accessible and restore the application services and data to users.

Figure 13-3 Recovery situation after a disaster occurs



Components of VVR that enable disaster recovery

This topic provides information about components of VVR that make the disaster recovery solution work.

Understanding replication

The term “replication” generally refers to the use of a tool or service, or a combination of tools or services, to automate the process of regularly placing an up-to-date copy of data from a designated source, or primary, to one or more remote locations.

Replication can be used to provide solutions to problems in a variety of application environments. Any application that needs redundancy at multiple sites or can achieve better performance through geographic distribution can benefit from replication. Redundancy at multiple sites, where updates to the primary site are immediately reflected at remote sites, can be effectively used to manage disaster recovery with the use of a replication tool.

Veritas Volume Replicator (VVR) is a data replication service that helps to maintain a consistent copy of the application data at a remote site. It is built to contribute to an effective disaster recovery plan. If the primary data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site. VVR works as an integrated component of Veritas Storage Foundation for Windows. Any application, even with existing data, can be configured to use VVR transparently. For more information on VVR, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

Modes of replication

VVR replicates in synchronous, asynchronous, and synchronous override modes.

Synchronous replication

The synchronous mode ensures that an update has been acknowledged by the secondary host before completing the update at the primary site. Thus, the primary site and the secondary site have the same data. If a disaster occurs on the primary site and its data is destroyed, the secondary site will already have an up-to-date copy of the data.

The synchronous mode of replication is most effective in application environments that have lower update rates but require all the hosts to always reflect the same data, or where a delay in updates between the primary and secondary hosts is not acceptable.

Asynchronous replication

In the asynchronous mode of replication, the application updates are immediately reflected at the primary site and sent to the secondary site as soon as possible. The updates are stored in the Replicator Log until they are sent to the secondary site. This allows asynchronous replication to deal with temporary network or secondary host failures without affecting the performance of the application.

Asynchronous replication mode is most effective in application environments where it is not acceptable for the application performance to be impacted, only a minimal data loss can be tolerated, or the application has a high rate of updates.

Synchronous override replication

The synchronous override mode of replication provides synchronous replication, as long as the network is available. If the network becomes unavailable, replication is continued in asynchronous mode.

The synchronous override replication mode is most effective in application environments where it is not acceptable for the primary site to be affected by a network failure.

Note: For additional information, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

Features of VVR that help in disaster recovery

While many of the components described above are replicated at the disaster recovery site through conventional means, VVR solves the difficult problem of replicating the *user* database. Refer to the following information on how VVR helps with disaster recovery in any application environment:

- **Write Order Fidelity:** VVR guarantees that changes made to data on the primary host are made in the same sequence on the secondary host. This ensures that the data remains in a consistent state in the event of a disaster.
- **Synchronous Replication:** VVR guarantees that changes committed on the primary host are committed on the secondary host first. This ensures that the data on the secondary host matches the data on the primary host and minimizes data loss in the event of a disaster.
- **Asynchronous Replication:** VVR reflects the changes to the application immediately on the primary, and changes are then reflected on the secondary as soon as possible. Until the data is sent to the secondary, it is stored on the Replicator Log.

- **RVG Snapshot:** This provides the ability within VVR to take a point-in-time snapshot of a volume. This allows verification of the consistency of the data on the secondary host without impacting replication between the primary and secondary hosts.
- **Heterogeneous Storage Support:** VVR provides a replication technology that works with heterogeneous storage hardware. VVR allows replication to occur between similar or dissimilar storage arrays from a vendor or between different storage arrays from different vendors. This allows for maximum use of existing hardware and provides flexibility when adding new hardware.

Components of VVR that enable disaster recovery

Deploying disaster recovery: New application installation

This chapter covers the following topics:

- [“Tasks for a new disaster recovery installation—additional applications”](#) on page 482
- [“Reviewing the requirements”](#) on page 485
- [“Reviewing the configuration”](#) on page 492
- [“Configuring the storage hardware and network”](#) on page 494
- [“Managing disk groups and volumes”](#) on page 497
- [“Setting up the secondary site: Configuring SFW HA and setting up a cluster”](#) on page 498
- [“Verifying that your application or server role is configured for HA at the primary site”](#) on page 520
- [“Setting up your replication environment”](#) on page 520
- [“Assigning user privileges \(secure clusters only\)”](#) on page 528
- [“Configuring disaster recovery with the DR wizard”](#) on page 529
- [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 533
- [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 537
- [“Installing and configuring the application or server role \(secondary site\)”](#) on page 541

- [“Cloning the service group configuration from the primary to the secondary site”](#) on page 543
- [“Configuring replication and global clustering”](#) on page 546
- [“Verifying the disaster recovery configuration”](#) on page 563
- [“Establishing secure communication within the global cluster \(optional\)”](#) on page 565
- [“Adding multiple DR sites \(optional\)”](#) on page 567
- [“Possible task after creating the DR environment: Adding a new failover node to a VVR environment”](#) on page 568
- [“Maintaining: Normal operations and recovery procedures \(VVR environment\)”](#) on page 571
- [“Recovery procedures for service group dependencies”](#) on page 574

Tasks for a new disaster recovery installation—additional applications

Before setting up disaster recovery at the secondary site, you must complete the high availability configuration for the application on the primary site.

See [Chapter 8, “Deploying SFW HA for high availability: New installation”](#) on page 77.

You can also configure disaster recovery for a primary site that is configured as a replicated data cluster.

See [Chapter 12, “Deploying Replicated Data Clusters: New application installation”](#) on page 333.

After setting up an SFW HA environment on the primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage configuration and the service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See “[Using the Solutions Configuration Center](#)” on page 29.

Note: If you want to create the identical configuration manually, without using the DR wizard, see [Appendix A, “Deploying Disaster Recovery: Manual implementation”](#) on page 873.

This chapter describes the process for any generic application or applications such as FileShare, PrintShare, IIS and MSVirtual Machines.

For examples of the SFW HA disaster recovery solution with specific applications, see the other Solutions Guides included with this release:

- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*
- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft Exchange.*

[Table 14-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 14-1 Task list for deploying disaster recovery

Objective	Tasks
“Reviewing the requirements” on page 485	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 492	Understanding active-passive configuration and site failover in a DR environment
“Configuring the storage hardware and network” on page 494	<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 the application will be installed
“Setting up the secondary site: Configuring SFW HA and setting up a cluster” on page 498	<ul style="list-style-type: none">■ Installing SFW HA■ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
“Verifying that your application or server role is configured for HA at the primary site” on page 520	Verifying that the application has been configured for high availability at the primary site

Objective	Tasks
“Setting up your replication environment” on page 520	Ensuring replication prerequisites for your selected method of replication are met before running the DR wizard
“Assigning user privileges (secure clusters only)” on page 528	For secure clusters only, assigning user privileges
“Configuring disaster recovery with the DR wizard” on page 529	<ul style="list-style-type: none"> ■ Reviewing prerequisites for the DR wizard ■ Starting the DR wizard and selecting a primary site system, the service group, the secondary site system, and the replication method
“Cloning the storage on the secondary site using the DR wizard (VVR replication option)” on page 533	(VVR replication option) Cloning the storage configuration on the secondary
“Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 537	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Using the DR wizard to create temporary storage for installation on the secondary site
“Installing and configuring the application or server role (secondary site)” on page 541	<ul style="list-style-type: none"> ■ Reviewing the prerequisite checklist ■ Installing the application
“Cloning the service group configuration from the primary to the secondary site” on page 543	Cloning the service group configuration from the primary to the secondary site using the DR wizard
“Configuring replication and global clustering” on page 546	<ul style="list-style-type: none"> ■ (VVR replication) Using the wizard to configure replication and global clustering ■ (EMC SRDF replication) Setting up replication and then using the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Setting up replication and then using the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Using the wizard to configure global clustering, and then setting up replication
“Verifying the disaster recovery configuration” on page 563	Verifying that the secondary site has been fully configured for disaster recovery

Objective	Tasks
“Establishing secure communication within the global cluster (optional)” on page 565	Adding secure communication between local clusters within the global cluster (optional task)
“Adding multiple DR sites (optional)” on page 567	Optionally, adding additional DR sites to a VVR environment
“Possible task after creating the DR environment: Adding a new failover node to a VVR environment” on page 568	Completing required tasks when adding a new failover system to either the primary or secondary site in a VVR environment
“Maintaining: Normal operations and recovery procedures (VVR environment)” on page 571	<ul style="list-style-type: none"> ■ Monitor replication ■ Perform planned migration ■ Complete the recovery procedures after the primary site goes down

Reviewing the requirements

Review the following installation and configuration requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

- “Disk space requirements” on page 485
- “Requirements for Veritas Storage Foundation High Availability for Windows (SFV HA)” on page 486

Review additional configuration topics before setting up disaster recovery.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. Table 14-2 estimates disk space requirements for SFV HA.

Table 14-2 Disk space requirements

Installation options	Install directory/drive
SFV HA + all options + client components	1564 MB
SFV HA + all options	1197 MB
Client components	528 MB

Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)

Before installing Veritas Storage Foundation High Availability for Windows (SFW HA), ensure that you review the following:

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 5.1 Service Pack 2 Hardware Compatibility List to confirm supported hardware:
<http://entsupport.symantec.com/docs/358407>
- Review the SFW HA 5.1 Service Pack 2 Software Compatibility List to confirm supported software:
<http://entsupport.symantec.com/docs/358406>
- Review the operating systems supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

Supported operating systems

SFW and SFW HA have client and server components that run on specific Windows operating systems.

SFW and SFW HA software for servers

Your system must run one of the operating systems listed in the tables to install the SFW or SFW HA server software:

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 14-3 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

The following table shows supported Windows Server 2008 operating systems.

Note: SFW HA supports Windows Server 2008 R2 without Hyper-V x64: Standard Edition, Enterprise Edition, Datacenter Edition.

Table 14-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			

Table 14-4 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server Note: SFW HA does not support Hyper-V.	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition Note: SFW HA does not support Hyper-V.	
Windows Server 2008 R2 for IA systems	IA64		
Windows Storage Server 2008			

SFW and SFW HA software for clients

The SFW or SFW HA client software can be installed on any one of the operating system versions, editions, and architectures that the Server Components are supported on, except Server Core.

The following table shows the additional client operating systems that are supported.

Table 14-5 Supported client operating systems

Windows Client	Platform	Edition	Service Pack
Windows XP	x86, x64		SP3 required
Windows Vista	x86, x64,	Ultimate Edition, Business Edition, Premium Edition	SP1 or SP2 required
Windows 7	x86, x64	Ultimate Edition, Business Edition, Premium Edition	

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs. See "[Best practices for SFW HA](#)" on page 492.
- NIC teaming is not supported for the VCS private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each application virtual server.
 - A minimum of one static IP address for each physical node in the cluster.

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements for SFW HA

The following permissions are required:

- You must be a domain user.
- You must be a member of the local Administrators group on all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements for SFW HA

Please review the following additional requirements:

- Installation media for all products and third-party applications.
- Licenses for all products and third-party applications.

- You must install the operating system in the same path on all systems. For example, if you install Windows Server on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- For a Replicated Data Cluster, install only in a single domain.

Best practices for SFW HA

Symantec recommends that you perform the following tasks:

- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the vxclus UseSystemBus ON command. This is applicable for a Replicated Data Cluster configuration.

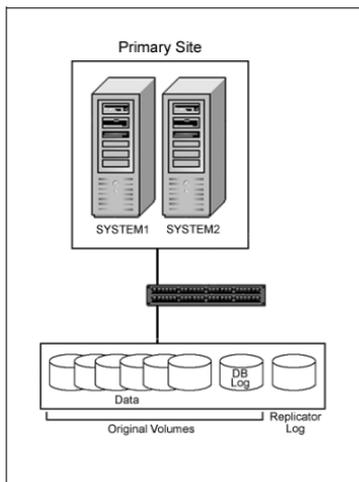
Reviewing the configuration

This configuration overview describes active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), then SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the DR

configuration that includes both sites, see the illustration in the section “[About a disaster recovery solution](#)” on page 474.

Figure 14-1 DR configuration primary site



Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported. If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

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.

Note: Follow the appropriate procedure for verifying the DNS settings, as the procedures for Windows Server 2003 and Windows Server 2008 are slightly different.

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 VCS 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.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 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.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 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 and Dial-up 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 VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in 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**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 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 VCS 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.
- 13 Click **OK**.

Managing disk groups and volumes

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - Control Panel - System Settings). See the *Veritas Storage Foundation Administrator's Guide* for more information.

This section includes the following procedures:

- Importing a disk group and mounting a shared volume
- Unmounting a volume and deporting a disk group

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**.

Setting up the secondary site: Configuring SFW HA and setting up a cluster

Begin with verifying that the requirements are met on the secondary site:

- See [“Reviewing the requirements”](#) on page 485.

Then continue with the following topics:

- [“Installing SFW HA”](#) on page 499
- [“Configuring the cluster”](#) on page 233

Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

When installing Veritas Storage Foundation HA for Windows, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 14-6 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.

Note: Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 500.

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.

Installing Symantec Trusted certificate for unsigned drivers

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 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.

If you select this option when installing on Windows Server 2003, you do not need to set the driver signing options to Warn or Ignore.

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

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 HA, 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, select the appropriate SFW product options, and click **Next**.

Client Components	Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used while configuring high availability.
-------------------	--

	Required to install the Solutions Configuration Center which provides information and wizards to assist configuration
--	---

Global Cluster Option	Required for a disaster recovery configuration only.
-----------------------	--

Volume Replicator (VVR)	If you plan to use Veritas Volume Replicator (VVR) for replication, select this option to install VVR.
-------------------------	--

High Availability Hardware Replication Agents	If you plan to use hardware replication, select the appropriate hardware replication agent.
---	---

- 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.

If you are using multiple paths and selected 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 path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.

If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:

On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).

When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:

When installing Veritas Storage Foundation for Windows (Server Components) with the Microsoft Cluster Server (MSCS) option, 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. For more information, see the *Veritas Storage Foundation Administrator's Guide*.

Review the messages, and then click **OK**.

- 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.

Resetting the driver signing options

After completing the installation sequence, reset the driver signing options on each computer.

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 the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management

Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
 - When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
 - Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

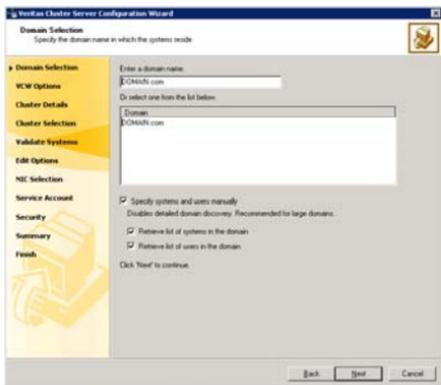
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



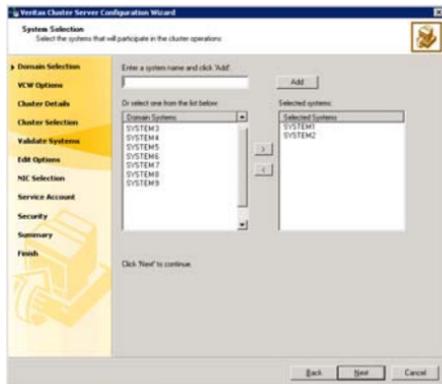
Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 507.

- To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**. If you chose to retrieve the list of systems, proceed to [step 6](#) on page 507. Otherwise, proceed to the next step.
- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster. Proceed to [step 8](#) on page 507.

- 6 On the System Selection panel, specify the systems for the cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the > (right-arrow) button.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

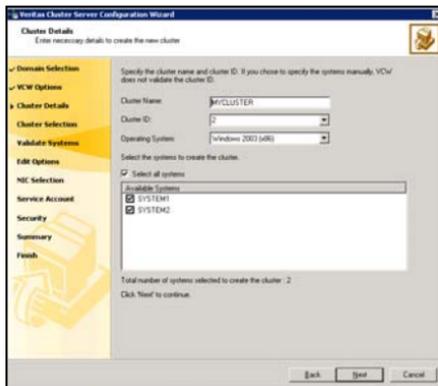
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

Operating System From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the network interface cards (NICs) on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

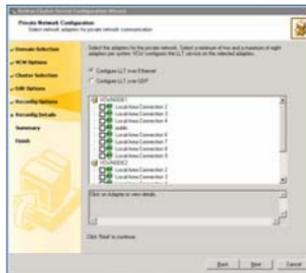
- 10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 511.

- 11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

- To configure the VCS private network over the ethernet, complete the following steps:



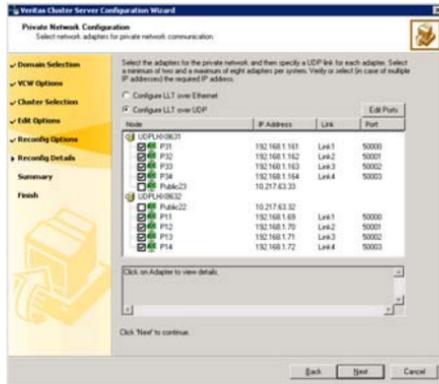
- Select the check boxes next to the two NICs to be assigned to the private network.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links.

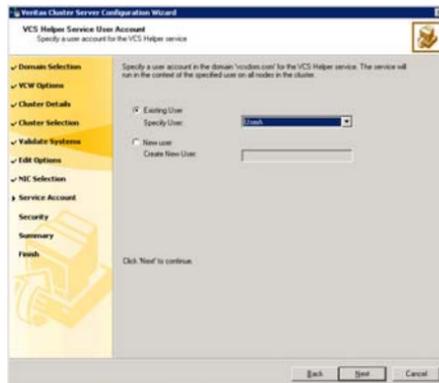
Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

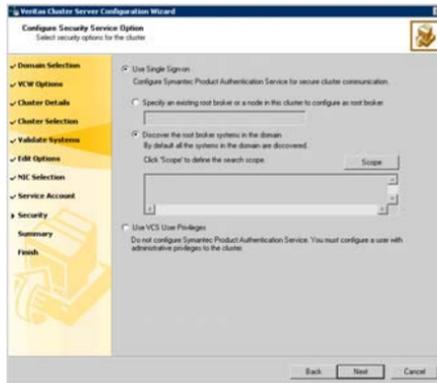
- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 506, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:
- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.
If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.
- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.
For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table 14-7 contains some more examples of search criteria.

Table 14-7 Search criteria examples

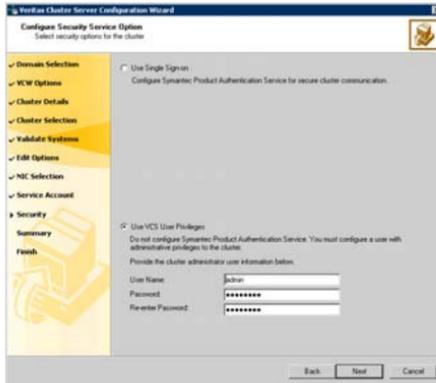
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
 - Click **Next**.
- 14 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

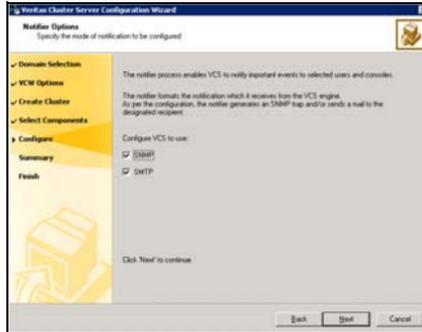
- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See "[Configuring notification](#)" on page 516.

Configuring notification

This section describes steps to configure notification.

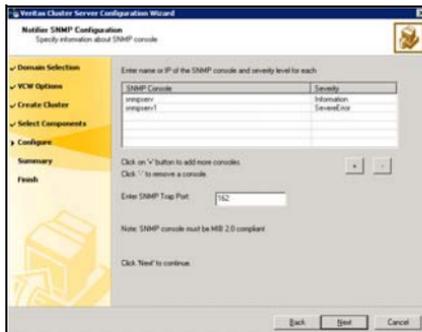
To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

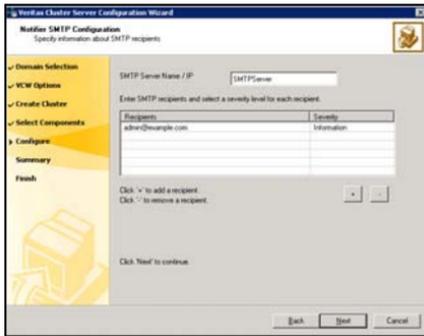


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

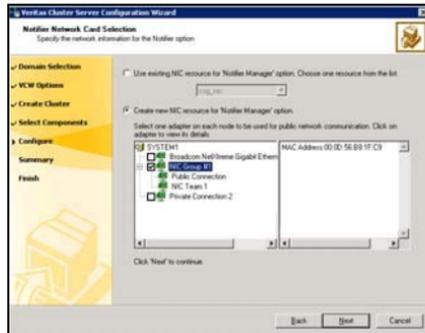


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.
 - Enter an SNMP trap port. The default value is "162".
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.
 - 6 Click **Configure**.
 - 7 Click **Finish** to exit the wizard.

Verifying that your application or server role is configured for HA at the primary site

Make sure that your application has been configured for high availability at the primary site. If you have not yet configured the application for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [Chapter 8, “Deploying SFW HA for high availability: New installation”](#) on page 77.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center. See [Chapter 12, “Deploying Replicated Data Clusters: New application installation”](#) on page 333.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Veritas Volume Replicator (VVR)
- EMC SRDF
- Hitachi TrueCopy

For array-based hardware replication, you can use any replication agent supported by Veritas Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only”](#) on page 560.

Before configuring replication with the wizard, ensure that you set up the replication environment prerequisites. Choose from the following topics, depending on which replication method you are using:

- [“Configuring the VVR security service”](#) on page 521

- “Requirements for EMC SRDF array-based hardware replication” on page 523
- “Requirements for Hitachi TrueCopy array-based hardware replication” on page 526

Configuring the VVR security service

If you are using Veritas Volume Replicator (VVR) replication, you must configure the VxSAS service on all cluster nodes on both the primary and secondary sites.

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server environment, 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 Launch the VVR Security Service Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxsascfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name (domain\account)	Enter the administrative account name.
----------------------------------	--

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

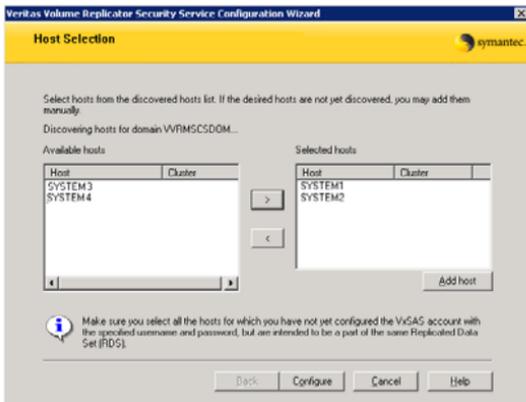
- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts	<p>The Available hosts pane lists the hosts that are present in the specified domain.</p> <p>Move the appropriate host 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.</p>
Adding a host	<p>If the host name you require is not displayed, click Add host. 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.</p>

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- Click **Finish** to exit the wizard.

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for configuring EMC SRDF”](#) on page 523
- [“Replication requirements for EMC SRDF”](#) on page 524

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all

Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
 - On the secondary site, the wizard verifies that no devices are RDF1.
- Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

In addition, note the following agent requirement:

- Device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

Before using the DR wizard, review the following topics:

- “Software requirements for Hitachi TrueCopy” on page 526
- “Replication requirements for Hitachi TrueCopy” on page 526

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

For details, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.
- Enter the primary and secondary site file paths for the horcm files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:
`System Driver\Windows`
- The horcm files are named `horcmnn.conf` (where *nn* is a positive number without a leading zero, for example, `horcm1.conf` but not `horcm01.conf`).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

In order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the application service group as well as any dependent service groups except for the RVG service group.

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

To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Modify the attribute of the service group to add the user. Specify the application service group and any dependent service groups except for the RVG service group.
`hauser -add user [-priv <Administrator|Operator> [-group service_groups]]`
- 4 Reset the configuration to read-only:
`haconf -dump -makero`

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Reset the configuration to read-only:
`haconf -dump -makero`

Configuring disaster recovery with the DR wizard

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (VVR replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure VVR replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

The wizard allows you to exit after the logical completion of each task. Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking **Next** then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.
- The clusters taking part in the DR configuration should have distinct names.

- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- One static IP address is available per application service group to be cloned.
- If using VVR for replication, a minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- [“Configuring the VVR security service”](#) on page 521
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 523
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 526

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

Note: By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

- 2 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.

3 In the System Selection panel, complete the requested information:

System Name	Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the application is online. If you have launched the wizard on the system where the application is online at the primary site, you can also specify <code>localhost</code> to connect to the system.
-------------	---

Click **Next**.

4 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.

You can choose to clone only the parent service group by not selecting the dependent service group. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.

Click **Next**.

6 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)	Select this option if you want to configure VVR replication. Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment. The wizard verifies each configuration task and recognizes if a task has been completed successfully. You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.
---	---

Configuring disaster recovery with the DR wizard

Configure EMC SRDF and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.</p> <p>Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.</p> <p>Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure the Global Cluster Option (GCO) only	<p>If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option.</p> <p>Select this option if you want to use the wizard in an array-based replication environment that is not supported by this wizard. You must configure replication manually after you finish the wizard.</p> <p>If you select the GCO only option, the DR wizard sets up the storage and service group configuration on the secondary site for an array-based hardware replication environment. Therefore, you cannot use this option to clone the storage and service group for a VVR replication environment.</p>

Click **Next**.

- 7 Continue with the next DR configuration task.
 - For VVR replication, see [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 533.
 - For array-based replication, see [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 537.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic before continuing with the storage cloning procedure:

[“Configuring disaster recovery with the DR wizard”](#) on page 529.

To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.
- 2 Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.

Recommended Action	<p>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.</p> <ul style="list-style-type: none">■ If the volume does not exist, a new volume will be created.■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.
--------------------	---

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site. Click **Next**.

- 3 In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the

primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.

Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4 In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group	Displays the disk group name to which the volume belongs.
Volume (Volume Size)	Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.
Available Disks	Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> option to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.
	Select disks for each unavailable volume that you want to clone on to the secondary.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.
View Primary Layout	Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.

- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the Application Installation panel, review the information and do one of the following:
 - If the application is already installed on the required nodes, click **Next** to continue with service group cloning. However, for Enterprise Vault, refer to the instructions in the documentation for configuring the Enterprise Vault service group manually and configuring Enterprise Vault for the cluster environment.
 - If the application is not yet installed on the secondary site, proceed with installation on the required nodes as follows:

For applications that require installing components on shared storage, before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.

click **Finish** to exit the wizard before proceeding with installation on that node. Afterwards, restart the Disaster Recovery wizard and continue through the wizard from the Welcome panel.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

To enable you to install applications, the DR wizard can create a temporary disk group, DR_APP_INSTALL__DG, which contains the volumes and mount points for use in application installation. The temporary configuration uses 500 MB volumes or the volume size at the primary site, depending on which is smaller. The wizard deletes the temporary configuration after application installation.

If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning check box on the Storage Cloning panel.

If you are starting the wizard for the first time, see the following topic before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 529.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface Default: C:\HORCM\etc where C is the system drive.
-----------------------	--

HORCM files location	Path to the horcm configuration files (horcm nn .conf) Default: C:\Windows where C is the system drive An horcm configuration file is required by the RAID Manager on all nodes; however the wizard does not validate this.
----------------------	--

- In the Storage Cloning panel, you can choose whether or not to perform storage cloning, which creates a temporary storage disk group and volumes for application installation. The wizard will delete the temporary storage once you confirm application installation is complete.

Choose one of the following:

- If you have not yet installed the application on all nodes, leave **Perform storage cloning** checked and click **Next**. Continue with the next step in this procedure.
 - If you have already installed the application on all nodes, uncheck **Perform storage cloning** and click **Next**. Continue with the procedure for service group cloning.
- The Storage Validation Results panel shows the temporary storage configuration that the wizard will configure at the secondary site. You can click **Show Summary** to toggle to a summary view and toggle back to a detailed view by clicking **Show Details**.

The detailed view shows the following:

Disk Group	Displays the name of the single disk group required on the secondary site for temporary storage: DR_APP_INSTALL_DG
Volume	Displays the list of volumes required at the secondary site.
Size	Displays the size of the volumes required on the secondary site.
Mount	Displays the mounts required at the secondary site.
Recommended Action	Indicates the action that the wizard will take at the secondary site.

The summary view shows the following.:

Existing configuration	Displays the existing secondary configuration.
------------------------	--

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.
---------------------------------	---

If the panel displays a message indicating that the available disks on the secondary are inadequate, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** so that the wizard can update its information about the secondary storage configuration.

Click **Next**.

- In the Disk Selection for Storage Cloning panel, a default disk selection is shown for the temporary storage at the secondary site. You can change the selection by moving disks to and from the Available Disks and Selected Disks pane. Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures. Click **Next**.
- The Volume Layout for Secondary Site Storage panel shows a default volume layout for the temporary storage based on the primary site volume layout. Optionally, you can change the default disk assignment and layout for any volume:

Disk Group	Displays the DR_APP_INSTALL__DG disk group.
Volume (Volume Size)	Displays the name and the size of the volume to be created on the secondary.
Available Disks	Displays the disks that are available for the volumes. To select a disk, either double-click on the host name or click the >> button to move the hosts into the Selected Disks pane.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been selected for the volume. To remove a disk from the list, select it and click the << button.
View Primary Layout	Displays the volume layout at the primary site.

Click **Next**.

- In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the temporary storage configuration at the secondary site.

- 9 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the Application Installation panel, review the information and do one of the following:
 - For applications that require installing components on shared storage, before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - If you keep the wizard running during installation, once application installation is complete, click **Next** to proceed with service group cloning. Otherwise, restart the DR wizard and continue through the wizard from the Welcome panel.

Once the application is installed, the temporary cloned storage is no longer needed. Before beginning service group cloning, the wizard displays the Temporary Storage Deletion panel to confirm the deletion of the temporary storage.

Installing and configuring the application or server role (secondary site)

This section provides important points which you must consider before you install the application.

The FileShare and PrintShare applications are installed as a part of the SFW HA installation. For any other application refer to the documentation provided with the application to complete the installation.

Installing the FileShare application

Points to note when installing FileShare:

- Make sure the disk group and volumes that contain the file server shared directory exist on the shared storage.
- When installing and configuring a new file server shared directory, create the disk groups and volumes on the shared storage and subsequently create the directory structure for the file shares on the shared storage.
- If your configuration already has a file server with shares on the local storage, then move these shares to the shared storage using practices recommended by Microsoft.

Installing the PrintShare application

Points to note when installing PrintShare:

- Make sure the required printer drivers have been added to all the systems in the cluster that are intended to be a part of the print share service group.
- For details, see *Veritas Cluster Server Administrator's Guide*.
- Make sure the printer is connected to the network and is configured with an IP address.

Installing the IIS application

Points to note when installing IIS:

- Verify IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on shared storage.
- Import the cluster disk groups and mount the volumes that contain the website data, on the first node.
- For a new IIS installation, while creating new websites, create the site folder on the shared storage and place the site content in that folder.

- Change the default home directory path for all IIS sites to be monitored to a location on the shared storage. See the IIS documentation for instructions.
- For existing websites, stop the sites and then move the website content to volumes on the shared storage. You must also reconfigure the home directory location for the website in IIS and then restart the website again.
- Verify the port numbers assigned to IIS sites are not used by other applications or sites.
- Synchronize the IIS configuration on all nodes hosting the service group.

Installing the Microsoft Virtual Machine application

Points to note when installing MS Virtual Machine:

- Verify Microsoft Virtual Server is installed and configured identically on all nodes hosting the service group.
- Install the operating system and the applications that you want to make highly available on the virtual machine.
- Install and configure Virtual Machine Additions *on each virtual machine* if you plan to enable detailed monitoring for the virtual machine resources.
- Verify the Microsoft Virtual Server configuration files reside locally on each node.
- Make sure the name of the virtual machine is unique in the cluster.

Installing additional applications

Following are some very generic points for installing any application:

- Make sure that the disk groups and volumes are mounted on the node. The DR wizard mounts volumes assigned to drive letters on the first node. However, if the primary site uses folder mounts, you must format the volumes and specify the drive path manually using the Veritas Enterprise Administrator. For each additional node, you must unmount volumes and deport the disk groups and import and mount them on the additional node.
- VCS requires the application program files to be installed on the same local drive on all nodes. For example, if you install the application program files on drive C of one node, installation of these same files on all other nodes must be on drive C. Make sure that the same drive letter is available on all nodes and has adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the shared storage.

Cloning the service group configuration from the primary to the secondary site

Before cloning the service group on the secondary site, verify that you have installed the application on the secondary site.

If you are launching the wizard for the first time, see the following topic for additional information:

[“Configuring disaster recovery with the DR wizard”](#) on page 529

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.
If you selected the VVR replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel.
If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.
- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:
 - If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.
 - If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.

- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.
- 6 Review the following information displayed in the Service Group Analysis panel and click **Next** to continue with service group cloning.

Service Group Name Displays the list of application-related service groups present on the cluster at the primary site.

Service Group Details on the Primary Cluster Displays the resource attributes for the service group at the primary site.
The NIC resource consists of the MAC address.
The IP resource consists of the IP address and subnet mask.

Service Group Details on the Secondary Cluster Displays a message to indicate whether the service group or the corresponding attributes have been configured at the secondary site.

- 7 In the Service Group Cloning panel, specify the requested system information for the secondary site.

Service Group Name Depending on the application service group already created at the primary site, and subsequently selected on the Service Group Selection page, the wizard displays the names of the service groups that will be cloned at the secondary site.

Available Systems Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning.
Select any additional secondary systems on which you want the wizard to clone the application service group configuration.
Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane.
Note: If you want to add systems to a service group after you finish cloning the service group configuration with the DR wizard, you cannot do so by running the DR wizard again. Instead, run the VCS configuration wizard and edit the system list of the existing service group.

Selected Systems Displays the list of selected systems. The secondary system that you selected earlier in the wizard is listed by default.

Click **Next**.

- 8 In the Service Group Attribute Selection panel, complete the requested information to create the required resources on the secondary site. The panel also displays the service group resource name and the attribute information at the primary site.

Resource Name	Displays the list of resources that exist on the primary cluster.
Attribute Name	Displays the attribute name associated with each of the resources displayed in the Resource Name column. If you need to edit additional attributes that are not shown, you must edit them manually on the secondary site service group once service group cloning is complete.
Primary Cluster	Displays the primary attribute values for each of the displayed attributes.
Secondary Cluster	The default is the same as the primary cluster. The same virtual IP address can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment. For the MACAddress attribute select the appropriate public NIC from the drop-down list.

Click **Next**.

- 9 In the Service Group Summary, review the attribute information that will be cloned on to the secondary cluster. Click **Back** to change any of the secondary service group attributes. Otherwise, click **Next** to proceed with cloning the service group configuration on the secondary site.
- 10 In the Implementation panel, wait until all the tasks are completed. The progress bar indicates the status of the tasks. Successful tasks are marked with a check symbol. If some task could not be completed successfully, the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**
- 11 If the cloning failed, review the troubleshooting information. Otherwise, click **Next** to continue with the replication and GCO configuration, or with GCO only, depending on which option you selected.

Optionally, you can exit the wizard at this point and launch the wizard again later. When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. Click **Next** to continue to the replication and/or GCO configuration task.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- [“Configuring VVR replication and global clustering”](#) on page 546
- [“Configuring EMC SRDF replication and global clustering”](#) on page 554
- [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 557
- [“Configuring global clustering only”](#) on page 560

Configuring VVR replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure VVR replication and global clustering.

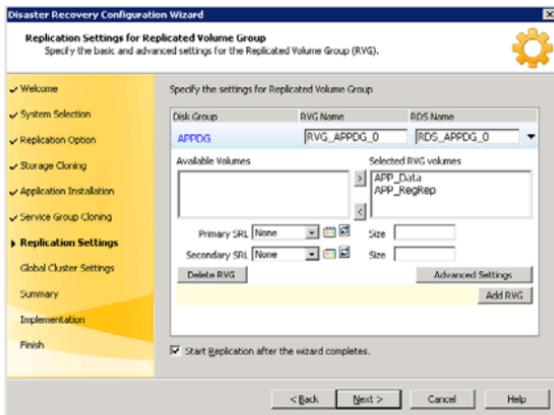
Before you begin, ensure that you have met the following prerequisites:

- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site. See the following topic:
 - [“Configuring the VVR security service”](#) on page 521
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that, for remote cluster configuration, you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure VVR replication and global clustering with the DR wizard.

To configure VVR replication and GCO

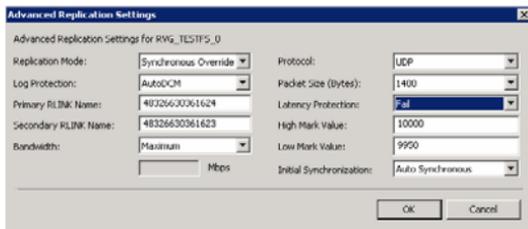
- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the wizard is still open after the previous wizard task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
- 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.



- | | |
|----------------------|--|
| Disk Group | The left column lists the disk groups. By design, an RVG is created for each disk group. |
| RVG Name | Displays the default RVG name. If required, change this to a name of your choice. |
| RDS Name | Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice. |
| Available Volumes | Displays the list of available volumes that have not been selected to be a part of the RVG.
Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane. |
| Selected RVG Volumes | Displays the list of volumes that have been selected to be a part of the RVG.
To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane. |

- | | |
|--|--|
| Primary SRL | <p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk.</p> <p>Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.</p> |
| Secondary SRL | <p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL.</p> <p>Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.</p> |
| Add RVG | <p>Click this option to create a new RVG. This option is especially useful if you want to organize the volumes present in a disk group under separate RVGs. By default the wizard is designed to organize all the volumes under a disk group under one RVG. However, using the Add RVG option you can choose to organize them differently, based on your specific requirements.</p> |
| Delete RVG | <p>Click this option to delete any of the existing RVGs related to the DR set up that you are creating.</p> |
| Start Replication after the wizard completes | <p>Select this check box to start replication automatically after the wizard completes the necessary configurations.</p> <p>Once replication is configured and running, deselecting the checkbox does not stop replication.</p> |
- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator*

Administrator's Guide for additional information on VVR replication options.



Replication Mode Select the required mode of replication: **Synchronous**, **Asynchronous**, or **Synchronous Override**. The default is synchronous override.

Log Protection Select the appropriate log protection from the list. 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.

The **Off** option disables Replicator Log Overflow protection.

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 **Fal** 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	Enter a name of your choice for the primary RLINK. If you do not specify any name then the wizard assigns a default name.
Secondary RLINK Name	Enter a name of your choice for the Secondary RLINK. If you do not specify any name then the wizard assigns a default name.
Bandwidth	<p>By default, VVR replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit.</p> <p>The default unit is Mega bits per second (Mbps) and the minimum allowed value is 1 Mbps.</p>
Protocol	Choose TCP or UDP. UDP/IP is the default replication protocol.
Packet Size (Bytes)	Default is 1400 Bytes. From the drop-down list, choose the required packet size for data transfer. The default unit for the packet size is Bytes. You can set the packet size only if the protocol is UDP/IP.
Latency Protection	<p>By default, latency protection is set to Off.</p> <p>When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection.</p> <p>This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.</p>
High Mark Value	<p>This option is enabled only when Latency Protection is set to Override or Fail. It specifies the maximum number of pending updates by which the secondary site can be behind the primary site. The default value is 10000.</p> <p>To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.</p>
Low Mark Value	This option is enabled only when Latency Protection is set to Override or Fail . When the updates in the Replicator log reach the High Mark Value , then the writes to the system at the primary site continues to be stalled until the number of pending updates on the Replicator log falls back to the Low Mark Value . The default is 9950.

Initial Synchronization	<p>If you are doing an initial setup, then use the Auto Synchronous option to synchronize the secondary site and start replication. This is the default.</p> <p>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.</p> <p>If you want to use the Synchronize from Checkpoint method then you must first create a checkpoint.</p> <p>If you have a 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 the checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.</p>
-------------------------	---

To apply changes to advanced settings, click **OK**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 5 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

Disk Group	Displays the list of disk groups that have been configured.
RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site.
Copy	Enables you to copy the above network settings to any additional RVGs that are listed on this screen. If there is only one RVG, this option does not apply.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings Allows you to use a WAC resource that already exists at either the primary or secondary site. Click **Primary** or **Secondary**, depending on the site at which the WAC resource already exists.

Resource Name Select the existing WAC resource name from the resource name list box.

Create new settings Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.

IP Address Enter a virtual IP for the WAC resource.

Subnet Mask Enter the subnet mask for the system at the primary site and the secondary site.

Public NIC Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.

Start GCO after configuration Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes.

Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.
Click **Next** to implement the settings.
- 8 In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description

about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have met the prerequisites for replication. See the following topic:

- “Requirements for EMC SRDF array-based hardware replication” on page 523

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings. See the following topic:

- “Optional settings for EMC SRDF” on page 556

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start** > **All Programs** > **Symantec** > **Veritas Cluster Server** > **Solutions Configuration Center**. Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration** > **Configure Disaster Recovery** > **Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.

- In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID) Specify the array ID for the primary site and for the secondary site.

Device Group name Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.

Available VMDG Resources Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings Allows you to use a WAC resource that already exists at either the primary or secondary site. Click **Primary** or **Secondary**, depending on the site at which the WAC resource already exists.

Resource Name Select the existing WAC resource name from the resource name list box.

Create new settings Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.

IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also detects and configures the SymHome attribute.

Other settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for EMC*

SRDF, Configuration Guide. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The optional settings use the following defaults:

Option	Default setting
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

Ensure that you have met the prerequisites. See the following topic:

- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 526

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings. See the following topic:

- [“Optional settings for HTC”](#) on page 560

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Additional

Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 4 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the horcm file is configured properly. If not, you can configure the horcm file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current horcm file.

- 5 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration,

GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

- 10 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only. For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1 If the wizard is still open after the service group cloning task, continue with the GCO Setup panel. Otherwise, launch the wizard and proceed to the GCO Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 2 In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.
- 3 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.

Start GCO after configuration

Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes.

Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 4 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. Click **Next**.
- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 6 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- For VVR replication:
 - Ensure VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.
 - Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
 - Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
 - Confirm that the RVG service groups are online at the primary and secondary sites.
 - Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.

Verifying the disaster recovery configuration

- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using VVR for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using VVR for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
 - starting a VVR replication checkpoint
 - performing a block level backup
 - ending the VVR replication checkpoint
 - restoring the block level backup at the DR site
 - starting replication from the VVR replication checkpointTo learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.
- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the `-secure` option to the `StartProgram` resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value.
For example:

```
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe"  
-secure
```
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.

Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel  
<low|medium|high> [--hashfile <filename> | --hash <root  
hash in hex>]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:

from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

from RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the `ClusterService-Proc (wac)` resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the **Process** agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

Possible task after creating the DR environment: Adding a new failover node to a VVR environment

The following procedure describes how to add an additional node to the cluster at either the primary or secondary site after your disaster recovery environment is in operation. The clusters at each site are not required to have the same number of nodes or the same failover configuration.

Preparing the new node

Install SFW HA on the new system and then add the system to the cluster.

To install SFW HA and add the system to the cluster

- 1 Refer to “[Installing SFW HA](#)” on page 499 for installation instructions.
- 2 Start the Veritas Cluster Server Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center** or click the shortcut for the **Solutions Configuration Center**.

From the **Solutions Configurations Center** expand **Disaster Recovery Configuration > Configure the cluster at the Secondary site** and from the display click **Configure the cluster** to add the new system to the cluster. If necessary, refer to the *Veritas Cluster Server Administrator’s Guide* for information on this procedure.

Preparing the existing DR environment

If you plan to add a failover node to the secondary site, you must temporarily switch the roles of the Primary and Secondary sites so that the current site becomes the Primary. This action reverses the direction of replication.

To prepare the existing DR environment

- 1 If you are adding the failover node to the cluster at the primary site, proceed directly to [step 2](#). If you are adding a failover node to the secondary site, you must switch the roles of the primary and secondary sites. This action reverses the direction of replication.
 - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
 - Click **Switch To**, and click **Remote switch**.
 - In the **Switch global group** dialog box, click the cluster at the secondary site you want to switch the group to. Then click the specific system

where you want to bring the global application service group online.
 Click **OK**.

- 2 Take the global application service group offline at the current primary site.
- 3 Take the VVR replication service group offline.

Modifying the replication and application service groups

Add the new failover node to the system lists in the replication and application service groups.

To add the failover node to the system lists

- 1 Bring the replication service group online on an existing cluster node of the current primary site.
- 2 Bring the MountV resources of the corresponding application service group online on the same node.
- 3 Use the **Modify an existing replication service group** option of the Volume Replicator Agent Configuration Wizard (**Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard**) to add a new node to the system list for the replication service group. If necessary, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide* for information on this procedure.
- 4 Use the **Modify service group** option of the FileShare, PrintShare, IIS, MSVirtual Server Machine Configuration Wizard or Application Configuration Wizard.
 Start the appropriate Configuration Wizard from the Solutions Configuration Center. For example, for FileShare click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center High Availability (HA) Configuration > Configure the Service Group > FileShare Configuration Wizard** to add the new node to the system list for the respective application service group. Check the check box to bring the service group online after the wizard completes. If necessary, refer to the *Cluster Server Administrator's Guide* for information on this procedure.
- 5 After bringing the application service group online, configure all the application database stores to automatically mount on start-up.

Reversing replication direction

If you added a failover node at the original secondary site and migrated the RVG in [“Preparing the existing DR environment”](#) on page 568, move the global application service group back to the original primary site and reverse the direction of replication. These actions switch the Primary and Secondary sites back to their original roles.

To reverse the replication direction

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group that is online at the current primary site.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the **Switch global group** dialog box:
 - Click the cluster to switch the group to.
 - Click the specific system where you want to bring the global application service group online.
 - Click **OK**.

Maintaining: Normal operations and recovery procedures (VVR environment)

This section provides tasks during normal operations of the DR solutions and also describes the recovery process.

Normal operations: Monitoring the status of the replication

Under normal operating conditions, you can monitor the status of the replication using the following tools:

- The VEA GUI
- The Command Line Interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Take the RVG resource offline on both the clusters.
- Transfer the primary role to the host at the secondary site by 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.
- Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- Bring the RVG resource online on the new secondary.
- Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

Note: Any changes that you make to the data on the new primary will get replicated to the original primary, which is now the secondary.

Disaster recovery procedures

This section provides information on bringing up an application server on the secondary host, in the event of a disaster. It also explains how to migrate the primary role back to the original primary host once it is returned to normal functioning after a disaster.

To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, right-click the desired secondary RVG node inside the replication network.
- 2 Select **Takeover** and follow the instructions to perform the takeover operation. You can choose to perform takeover with the following options:
 - Perform **Takeover** *with* the **fast-fallback** option to restore the original primary easily once it becomes available again. When performing **Takeover** with **fast-fallback**, make sure that you do not select the **Synchronize Automatically** option.
 - Perform **Takeover** *without* the **fast-fallback** option. In this case, you will need to perform a complete synchronization of the original primary with the new primary. This may take quite a while, depending on the size of the data volume. Only after the synchronization is complete can you migrate the primary role back to the original primary.
- 3 After the takeover, the existing secondary becomes the new primary.
- 3 Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online.

Now you can start using the application on the new primary.

Restoring the primary host

After a disaster, when the original primary becomes available again, you may want to revert the role of the primary back to this host.

To restore the primary host

- 1 Take the RVG resource off-line on both the clusters.
- 2 Depending on whether you performed **Takeover** *with* or *without* the **fast-failback** option, do one of the following:
 - For **Takeover** *with* the **Fast-failback** option:

The original primary, after it has recovered, will be in the **Acting as secondary** state. If the original primary is not in the **Acting as secondary** state, verify whether your network connection has been restored.

To synchronize this original primary and the new primary, use the **Resynchronize Secondaries** option from new primary's context menu.
 - For **Takeover** *without* the **Fast-failback** option:

After performing a takeover without fast-failback, you must convert the original primary to a secondary by 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 RDS's. 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 the **Synchronize Automatically** option.
- 3 After the synchronization is complete, perform a migrate operation to transfer the primary role back to the original primary. Right-click on the primary RVG and select **Migrate** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the RVG resource online on the secondary.
- 6 Bring the application group online on the original primary.

Recovery procedures for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

See “[Supported disaster recovery configurations for service group dependencies](#)” on page 493.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 14-8 Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none">■ The parent remains online on the primary site.■ An alert notification at the secondary site occurs for the child service group only.■ The RVG group remains online.	<ol style="list-style-type: none">1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).

Table 14-8 Online, local, soft dependency link

Failure condition	Results	Action required
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Table 14-9 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Recovery procedures for service group dependencies

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, hard dependency link.

Table 14-10 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Do not take the RVG group offline at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<p>1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.</p> <p>2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p>

Testing fault readiness by running a fire drill

Topics in this chapter include:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [About post-fire drill scripts](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Recreating a fire drill configuration that has changed](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)

About disaster recovery fire drills

A disaster recovery plan should include regular testing of an environment to ensure that a DR solution is effective and ready should disaster strike. This testing is called a fire drill.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group uses a copy of the data that is used by the application service group.

About the Fire Drill Wizard

Veritas Storage Foundation HA for Windows (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Veritas Volume Replicator (VVR) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

About Fire Drill Wizard general operations

The Fire Drill Wizard performs the following operations:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group. When creating the fire drill service group, the wizard uses the application service group name, with the suffix `_fd`. The wizard renames the fire drill service group resources with a prefix `FDnn` and changes attribute values as necessary to refer to the FD resources.
The wizard also supports fire drill service groups created under a different naming convention by an earlier version of the wizard.
- Runs the fire drill by bringing the fire drill service group online on the secondary site
This operation demonstrates the ability of the application service group to failover and come online at the secondary site should the need arise. Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.
- Restores the fire drill configuration, taking the fire drill service group offline
After you complete the fire drill, you run the wizard to restore the fire drill configuration to a prepared state. Otherwise, the fire drill service group remains online. If you run a fire drill on one service group, restore that service group before you continue with a fire drill on another service group. You must also restore the fire drill configuration before you can delete it.

Warning: If the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. Therefore, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible after completing the fire drill testing for a service group.

See [“Restoring the fire drill system to a prepared state”](#) on page 599.

- Deletes the fire drill configuration

The details of some Fire Drill Wizard operations are different depending on the replication environment.

See [“About Fire Drill Wizard operations in a VVR environment”](#) on page 579.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 580.

About Fire Drill Wizard operations in a VVR environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See [“About the Fire Drill Wizard”](#) on page 578.

However, the following additional Fire Drill Wizard operations are specific to a Veritas Volume Replicator (VVR) environment.

Preparing the fire drill configuration

In a VVR environment, when preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, replaces the RVGPrimary resources with VMDg resources
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill

You assign one or more disks for the mirrored volumes while running the wizard. Mirror preparation can take some time, so you can exit the wizard once this step is started and let the preparation continue in the background.

Running the fire drill

In a VVR environment, when running the fire drill, the wizard does the following:

- Detaches the mirrors from the original volumes to create point-in-time snapshots of the production data
- Creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes

Restoring the fire drill configuration

In a VVR environment, when restoring the fire drill system to a prepared state, the wizard does the following:

- Takes the fire drill service group offline
- Disables the fire drill service group resources
- Imports the fire drill disk group
- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

Deleting the fire drill configuration

In a VVR environment, when deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group and any associated registry entry
- Performs the snap abort operation on the snapshot mirrors to free up the disk space

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site
- Runs the fire drill, bringing the fire drill service group online on the secondary site

- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See “[About the Fire Drill Wizard](#)” on page 578.

However, additional Fire Drill Wizard operations are specific to a Hitachi TrueCopy or EMC SRDF replication environment.

In a Hitachi TrueCopy or EMC SRDF replication environment, the wizard performs the following additional actions during preparation, running of the fire drill, restoring the configuration, and deleting the configuration. You must configure the ShadowImage (for Hitachi) or BCV (for SRDF) pairs before running the wizard.

Preparing the fire drill configuration

When preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, the wizard creates HTCSnap or SRDFSnap resources for each HTC and SRDF resource in the application service group. The SRDFSnap and HTCSnap resources from the fire drill service group are linked to the respective resources configured in the main application service group.
- The wizard configures the Snap resource. The following Snap resource attributes are set to a value of 1:
 - UseSnapshot (take a local snapshot of the target array)
 - RequireSnapshot (require a successful snapshot for the Snap resource to come online)
 - MountSnapshot (use the snapshot to bring the fire drill service group online)

Running the fire drill

When running the fire drill, the wizard brings the HTCSnap or SRDFSnap agent online. The HTCSnap or SRDFSnap agent manage the replication and mirroring functionality according to the attribute settings. The Snap agents take a consistent snapshot of the replicating data using the mirroring technology provided by the array vendor. The Snap agents also import the disk group present on the snapshot devices with a different name.

In more detail, the Snap agent does the following:

- Suspends replication to get a consistent snapshot
- For HTCSnap, takes a snapshot of the replicating application data on a ShadowImage device

- For SRDFSnap, takes a snapshot of the replicating application data on a BCV device
- Resumes replication
- Modifies the disk group name in the snapshot

Restoring the fire drill configuration

When restoring the fire drill configuration to a prepared state, the wizard does the following:

- Takes the fire drill service group offline, thus also taking offline the SRDF and HTC Snap agents
This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

Deleting the fire drill configuration

When deleting the fire drill configuration, the wizard does the following:

- Deletes the fire drill service group
- Deletes any associated registry entry

If you want to remove the hardware mirrors, you must do so manually.

For more information about the Hitachi TrueCopy Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

About post-fire drill scripts

You can specify a script for the Fire Drill Wizard to run on the secondary site at the end of the fire drill.

For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

Note: The wizard does not support using script commands to launch a user interface window. In such a case, the process is created but the UI window does not display.

Optionally, you can specify to run a Windows PowerShell cmdlet by creating a .bat file.

To run a cmdlet, create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -command  
"$ScriptName"
```

Where

\$ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -command  
C:\myTest.ps1
```

Specify the name of the .bat file as the script to run.

Tasks for configuring and running fire drills

While running the Fire Drill Wizard, the following sequence of actions are available:

- Prepare the fire drill configuration
- Run the fire drill or delete the configuration
- Restore the fire drill configuration after running a fire drill
- Run another fire drill or delete the configuration

In addition, you have the option to recreate a fire drill configuration that has changed.

After an action is complete, the next action becomes available in the wizard. You can select the next action or exit the wizard and perform the next action later.

[Table 15-1](#) gives more details of the process of configuring and running fire drills with the wizard.

Table 15-1 Process for configuring and running fire drills

Action	Description
Verify the hardware and software prerequisites	Before running the wizard, review the prerequisites and make sure that they are met. See " Prerequisites for a fire drill " on page 586.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See " Preparing the fire drill configuration " on page 589.
Recreate a fire drill configuration that has changed	If a fire drill configuration exists for the selected service group, the wizard checks for differences between the fire drill service group and the application service group. If differences are found, the wizard can recreate the fire drill configuration before running the fire drill. See " Recreating a fire drill configuration that has changed " on page 597.

Table 15-1 Process for configuring and running fire drills

Action	Description
Run the fire drill	<p>Use the wizard to run the fire drill. Running the fire drill brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>See “Running a fire drill” on page 595.</p> <p>Perform your own tests of the application to confirm that it is operational</p> <p>Note: After completing the fire drill testing, run the wizard again as soon as possible to restore the configuration. Otherwise the fire drill service group remain online. Be sure to restore one fire drill service group to a prepared state before running a fire drill on another service group.</p>
Restore the fire drill configuration to a prepared state	<p>Use the wizard to restore the fire drill system to a state of readiness for future fire drills or to prepare for removal of the fire drill configuration</p> <p>This is a required action after running the fire drill.</p> <p>See “Restoring the fire drill system to a prepared state” on page 599.</p> <p>This operation reattaches snapshot mirrors and takes the fire drill service group offline.</p>
Delete the fire drill configuration	<p>If a fire drill service group is no longer needed, or if you want to free up resources, use the wizard to remove the fire drill configuration</p> <p>See “Deleting the fire drill configuration” on page 600.</p> <p>The wizard deletes the service group on the secondary site. In a VVR environment, the wizard performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill. In hardware replication environments, you can delete these manually.</p> <p>If a fire drill has been run, the wizard ensures that you first restore the fire drill configuration to a prepared state before this option becomes available. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p>

Prerequisites for a fire drill

Before running the Fire Drill Wizard make sure that you meet the following general requirements:

- You can run the Fire Drill Wizard from any node in the domain of the cluster, as long as the SFW HA client is installed on that node.
- If the cluster is secured, the login you use to run the Fire Drill Wizard must have the appropriate permissions to make changes in the cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall must be set to allow both ingoing and outgoing TCP requests on port 7419.
- For each IP address in the application service group, an IP address must be available to use on the secondary site for the fire drill service group.
The wizard can accept input for one IP address and Lanman resource. If the application service group has multiple IP addresses and Lanman resources, the wizard notifies you to edit the fire drill service group resources to supply these values. Information on editing service group resources is covered in the VCS administration guide.
See *Veritas Cluster Server Administrator's Guide*.
- If you want the fire drill wizard to run a script that you supply, ensure that the script file is available on any secondary site nodes where you plan to run the fire drill.

Note: The Fire Drill Wizard is not supported with a PrintShare service group.

Additional requirements apply to specific replication environments.

See ["Prerequisites for a fire drill in a VVR environment"](#) on page 586.

See ["Prerequisites for a fire drill in a Hitachi TrueCopy environment"](#) on page 587.

See ["Prerequisites for a fire drill in an EMC SRDF environment"](#) on page 588.

Prerequisites for a fire drill in a VVR environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See ["Prerequisites for a fire drill"](#) on page 586.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Veritas Volume Replicator (VVR) environment:

- The primary and secondary sites must be fully configured with VVR replication and the global cluster option.
- The Veritas FlashSnap option must be installed on all nodes of the secondary site cluster.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- On the secondary site, empty disks must be available with enough disk space to create snapshot mirrors of the volumes. Snapshot mirrors take up the same amount of space as the original volumes. In addition, two disk change object (DCO) volumes are created for each snapshot mirror, one for the source volume and one for the snapshot volume. The two DCO volumes must be on different disks. Allow 2 MB additional space for each DCO volume.
The empty disks must be in the same disk group that contains the RVG. If the disk group does not have empty disks available, you must use the VEA to add the disks to the disk group before you run the wizard. The secondary system must have access to the disks or LUNs.
- All disk groups in the service group must be configured for replication. The Fire Drill wizard does not support a VVR configuration in which disk groups are excluded from replication. However, you can exclude individual volumes within a disk group from replication.

Prerequisites for a fire drill in a Hitachi TrueCopy environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment. General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 586.

Make sure that the following prerequisites are met before configuring and running a fire drill in a Hitachi TrueCopy environment:

- The primary and secondary sites must be fully configured with Hitachi TrueCopy replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with Hitachi TrueCopy.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- Make sure that Hitachi RAID Manager/Command Control Interface (CCI) is installed.

- ShadowImage for TrueCopy must be installed and configured for each LUN on the secondary site target array. ShadowImage pairs must be created to allow for mirroring at the secondary site.
- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number should be different.
- Make sure the HORCM instance managing the S-VOLs runs continuously; the agent does not start this instance.

Prerequisites for a fire drill in an EMC SRDF environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment. General requirements are covered separately.

See ["Prerequisites for a fire drill"](#) on page 586.

Make sure that the following prerequisites are met before configuring and running a fire drill in an EMC SRDF environment:

- The primary and secondary sites must be fully configured with EMC SRDF replication and the global cluster option. The configuration must follow the applicable instructions in the Veritas Storage Foundation HA for Windows documentation for configuring disaster recovery with EMC SRDF.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. This process involves associating Symmetric Business Continuance Volumes (BCVs) and synchronizing them with the secondary site source.
- If you plan to run a fire drill on SRDF/A devices, you must have a TimeFinder/CG license. Make sure TimeFinder for SRDF is installed and configured at the target array.
- To take snapshots of R2 devices, BCVs must be associated with the RDF2 device group and fully established with the devices.
- To take snapshots of non-replicated devices, create a EMC Symmetrix device group with the same name as the SFW disk group. The device group must contain the same devices as in the disk group and have the corresponding BCVs associated.

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group. You specify the application service group and the secondary system to use. Only one service group can be prepared for a fire drill at one time.

For a Veritas Volume Replicator (VVR) environment, the preparation step also prepares snapshot mirrors of production data at the specified node on the secondary site.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration with the Fire Drill Wizard, make sure that you meet the prerequisites.

See “[Prerequisites for a fire drill](#)” on page 586.

To prepare the fire drill configuration

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
See “[System Selection panel details](#)” on page 591.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**.
See “[Service Group Selection panel details](#)” on page 591.
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.
See “[Secondary System Selection panel details](#)” on page 591.
- 7 If the Fire Drill Prerequisites panel is displayed, review the information and ensure that all prerequisites are met. Click **Next**.
See “[Prerequisites for a fire drill](#)” on page 586.

Otherwise, if a fire drill service group already exists on this system for the specified service group, one of the following panels is displayed:

If the Run Fire Drill option or Delete Fire Drill options are shown, a fire drill service group has already been prepared. You can run the fire drill with no further preparation. Click **Run Fire Drill** and follow the procedure for running a fire drill. See "[Running a fire drill](#)" on page 595.

If the Fire Drill Restoration panel is displayed, the fire drill service group remains online from a previous fire drill. Follow the procedure for restoring the fire drill configuration to a prepared state. This must be done before running a new fire drill. See "[Restoring the fire drill system to a prepared state](#)" on page 599.

If the Recreate Fire Drill Service Group panel is displayed, a fire drill service group has already been prepared but is not up to date. You can choose to recreate the fire drill configuration to bring it up to date. See "[Recreating a fire drill configuration that has changed](#)" on page 597. Or you can clear the check box to recreate the configuration and run the fire drill on the existing configuration.

- 8 If the Fire Drill Service Group Settings panel is displayed, assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. See "[Fire Drill Service Group Settings panel details](#)" on page 592.
- 9 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication Disk Selection panel
See "[Disk Selection panel details](#)" on page 592.

Hitachi TrueCopy replication Horcm Files Path Selection panel
See "[Hitachi TrueCopy Path Information panel details](#)" on page 593.
HTCSnap Resource Configuration panel
See "[HTCSnap Resource Configuration panel details](#)" on page 593.

EMC SRDF
replication

SRDFSnap Resource Configuration panel
See [“SRDFSnap Resource Configuration panel details”](#) on
page 594.

Click **Next**.

- 10 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.
See [“Fire Drill Preparation panel details”](#) on page 594.
When preparation is complete, click **Next**.
- 11 The Summary panel displays the message that preparation is complete.
To run the fire drill now, click **Next**. Continue with the procedure to run the fire drill.
See [“Running a fire drill”](#) on page 595.
To run the fire drill later, click **Finish**. The fire drill preparation remains in place.

System Selection panel details

Use the System Selection panel of the wizard to specify a system in the primary site cluster.

All systems containing online global service groups are available to select. The default system is the node where you launched the wizard (localhost) if a global service group is online on that system. When selecting a system you can specify either a fully qualified host name or IP address.

Service Group Selection panel details

Use the Service Group Selection panel of the wizard to select the service group that you want to use for the fire drill. You can select only one service group at a time for a fire drill.

Secondary System Selection panel details

Use the Secondary System Selection panel of the wizard to select the cluster and the system to be used for the fire drill at the secondary site.

The selected system must have access to the replicated data.

The system must have access to disks for the snapshots that will be created for the fire drill.

Fire Drill Service Group Settings panel details

Use the Fire Drill Service Group Settings panel of the wizard to assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. These must be an address and name not currently in use.

If the service group contains more than one IP and Lanman resource, this panel does not display. After the fire drill service group is created, the wizard notifies you to manually update the IP and Lanman resources in the fire drill service group.

Disk Selection panel details

During fire drill preparation in a VVR replication environment, you must ensure that information is available to the wizard for creating the fire drill snapshots. Use the Disk Selection panel of the wizard to review the information on disks and volumes and make the selections for the fire drill snapshots, as follows:

Volume	Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message. Note: The Disk Selection panel also appears if the wizard is recreating a fire drill service group to which volumes have been added. In that case, only the new volumes are shown for selection.
Disk Group	Shows the name of the disk group that contains the original volumes. This field is display only.

Fire Drill DG	Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with <i>FDnn</i> .
Disk	<p>Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume.</p> <p>You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group.</p> <p>If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.</p>
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This field is display only.

Hitachi TrueCopy Path Information panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the Hitachi TrueCopy Path Information panel is displayed.

The wizard populates the path field with the customary default location:

C:\Windows

where C is the system drive.

If the horcm configuration files are in a different location, edit the field to specify that location.

HTCSnap Resource Configuration panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the wizard discovers the HTC resources and non-replicating SFW disk groups in the application service group

This information is used to configure the HTCSnap resources.

The wizard lists each HTCSnap resource that will be configured. You can clear the HTCSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

You must specify the ShadowImage instance.

The HTCSnap Resource Configuration panel shows the following:

Target Resource Name	The panel shows the HTC resource name in the case of a Replication Device Group or the disk group resource name in the case of a non-replicating disk group.
ShadowImage Instance ID	For every HTC resource, specify the ID of the ShadowImage instance associated with the replicating secondary devices.
Refresh	If you click the Refresh button, the wizard rediscovers and validates the HTC configuration.

More information about HTCSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 580.

SRDFSnap Resource Configuration panel details

During fire drill preparation in an EMC SRDF replication environment, the wizard validates whether at least one BCV device is attached to every device (RDF2) of the SRDF device group. If not, the wizard displays an informational message on this panel. The panel shows as the Target Resource Name the name of the resource that is managing the LUNs that you want to snapshot. For data being replicated from the primary site, the Target Resource Name is the name of the SRDF resource. For data that is not replicated, the Target Resource Name is the name of the disk group resource.

The wizard lists each SRDFSnap resource that will be configured. You can clear the SRDFSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

More information about SRDFSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 580.

Fire Drill Preparation panel details

After you enter the information required to prepare a fire drill configuration, the Fire Drill Preparation panel is displayed. You wait while the wizard completes the preparation tasks.

The fire drill service group is created on the secondary site (but remains offline).

In addition, for a VVR replication environment, the snapshot mirrors for the volumes are prepared; this can take some time. You may want to minimize the wizard while the task runs in the background. You can also track the mirror preparation progress in the VEA. When done, the wizard displays a message that the fire drill preparation is complete. If the wizard is completing the preparation steps as part of recreating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See [“Recreating a fire drill configuration that has changed”](#) on page 597.

Running a fire drill

After you complete the initial fire drill preparation step using the Fire Drill Wizard, you can run the fire drill immediately without exiting the wizard or run the wizard later to run the fire drill.

Running the fire drill does the following:

- Creates the snapshots
- Enables the firedrill resources
- Brings the fire drill service group online
- Optionally, executes a specified command to run a script
See [“About post-fire drill scripts”](#) on page 582.

For details on the operations that occur when running a fire drill, see the following topics:

- [“About Fire Drill Wizard operations in a VVR environment”](#) on page 579
- [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 580

Warning: After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, run the wizard again to restore the system to the prepared state. Otherwise, if the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. See [“Restoring the fire drill system to a prepared state”](#) on page 599.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, or if you are returning to this procedure after recreating a fire drill service

group, go to [step 8](#). Otherwise, if you need to restart the wizard, continue with the next step.

- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
- 5 In the Service Group Selection panel, select the service group and click **Next**.
- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site and click **Next**.
If the fire drill configuration is in a prepared state, the wizard compares the resources of the fire drill service group with the resources of the application service group.
- 7 If the application service group changed since the fire drill configuration was prepared, the wizard displays the Recreate Fire Drill Service Group panel, showing the differences. Choose one of the following:
 - Leave the option checked to recreate the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.
For more information, see “[Recreating a fire drill configuration that has changed](#)” on page 597.
 - To run the fire drill on the existing configuration, clear the option to recreate the fire drill service group and click **Next**.
- 8 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**.
- 9 In the Post Fire Drill Script panel, optionally specify the full path to a script for the wizard to run on the secondary system right after running the fire drill. The script must already exist on the secondary system. Click **Next**.
See “[About post-fire drill scripts](#)” on page 582.
- 10 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and click **Next**. The Summary panel displays the message that the fire drill is complete. You can leave the wizard running while you verify the results or exit the wizard. To exit the wizard, click **Finish**.
- 11 Run your own tests to verify the fire drill results.

Warning: You should always restore the fire drill system to a prepared state immediately after completing fire drill testing on a service group.

- 12 Restore the fire drill configuration to the prepared state.
See “[Restoring the fire drill system to a prepared state](#)” on page 599.

Recreating a fire drill configuration that has changed

When you run the Fire Drill wizard, a fire drill service group may already exist for the selected application service group. However, the application service group may have changed since the fire drill service group was created. Therefore, the wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Recreate Fire Drill Service Group panel.

You have the following choices from the Recreate Fire Drill Service Group panel:

- Leave the option checked to recreate the fire drill service group.
Proceed with using the wizard to recreate the configuration to match the application service group.
The wizard deletes the existing fire drill configuration first, before creating the new one.
For a VVR replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it recreates the fire drill configuration, it prepares new mirrors only for new volumes. If volumes have been removed, the wizard displays an additional option to snap abort the obsolete snapshot volumes to free up disk space.
- Clear the option to recreate the fire drill service group. You can then proceed with using the wizard to do either of the following:
 - Run the fire drill, ignoring the differences.
 - Delete the entire fire drill configuration. Then start over with preparing the fire drill configuration.

Note: The wizard does not check for changes in volume attributes, such as the MountPath attribute. For example, if you have a MountV resource with an attribute that points to drive Y and you change that attribute to point to drive X, the wizard does not identify this change and does not give the option to recreate the fire drill service group.

You can choose whether to manually edit the fire drill service group for such changes and then run the fire drill, ignore the differences, or delete the configuration and start over.

The following procedure describes the choice of recreating the fire drill configuration.

To recreate the fire drill configuration if the service group has changed

- 1 In the Recreate Fire Drill Service Group panel, leave the option checked to recreate the configuration before running the fire drill.
For a VVR replication environment, if volumes have been removed, optionally select to snap abort the volumes.
Click **Next**.
- 2 In the Fire Drill Mode Selection panel, Delete Fire Drill Configuration is selected. Click **Next**, and click **Yes** to confirm the deletion.
- 3 The Fire Drill Deletion panel shows the progress of the deletion.
For a VVR replication environment, the wizard leaves the existing fire drill snapshot volumes so that those snapshot mirrors do not have to be prepared again. If volumes were removed and you selected the option to snap abort, the wizard snap aborts the snapshots of those volumes.

Warning: If you close the wizard after deleting the fire drill configuration without continuing on to the fire drill preparation step, the information of the existing snapshot volumes is lost.

When all tasks are complete, click **Next**.

- 4 In the Fire Drill Prerequisites panel, review the information and ensure that all prerequisites are met. Click **Next**.
See [“Prerequisites for a fire drill”](#) on page 586.
- 5 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

VVR replication	If volumes have been added, the Disk Selection panel is displayed. Specify the information for the added volumes. If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard. See “Disk Selection panel details” on page 592.
-----------------	--

Hitachi TrueCopy replication	Horcm Files Path Selection panel See “ Hitachi TrueCopy Path Information panel details ” on page 593. HTCSnap Resource Configuration panel See “ HTCSnap Resource Configuration panel details ” on page 593.
EMC SRDF replication	SRDFSnap Resource Configuration panel See “ SRDFSnap Resource Configuration panel details ” on page 594.

Click **Next**.

- The Fire Drill Preparation panel is displayed. Wait while the wizard recreates the fire drill service group.
For VVR replication environments, wait while the wizard starts mirror preparation.
Mirror creation can take some time. You may want to minimize the wizard while the task runs in the background. You can also close the wizard and track the mirror preparation progress in the VEA.
- Once preparation is complete, click **Next**. The Summary page is displayed. To continue with running the fire drill, click **Next**.
See “[Running a fire drill](#)” on page 595.

Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard as soon as possible to restore the fire drill system at the secondary site to a prepared state. A prepared state is the initial fire drill configuration created by the wizard, in which the fire drill service group has been prepared but is offline. Restoring the fire drill system to a prepared state is required for any of the following:

- Making the secondary system available for failover of the application service group at the primary site.
- Running another fire drill.
- Deleting the fire drill configuration after a fire drill has been run.

For details on the operations that occur when restoring a fire drill configuration, see the following topics:

- “[About Fire Drill Wizard operations in a VVR environment](#)” on page 579

- “About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment” on page 580

To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to [step 8](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Restoration Information panel, review the requirements for restoration and click **Next**.
- 8 In the Fire Drill Restoration screen, wait until the screen shows the restoration tasks are completed and click **Next**.
- 9 In the Summary screen, click **Next** if you want to delete the fire drill configuration. Otherwise click **Finish** to exit the wizard, leaving the fire drill configuration in a prepared state.

Deleting the fire drill configuration

If you no longer need a fire drill configuration you can delete it.

Deleting a fire drill configuration deletes the fire drill service group on the secondary site.

In a VVR replication environment, deleting a fire drill configuration also performs a snap abort of the snapshot mirrors created on the secondary site for use in the fire drill. It frees up the disk space used for the snapshot mirrors for other use.

In a Hitachi TrueCopy or EMC SRDF environment, you could manually remove mirrors after the deletion is complete.

To delete a fire drill configuration

- 1 If you have just used the wizard to prepare or restore a fire drill configuration and have not exited the wizard, go to [step 10](#). Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Additional Applications**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 If the wizard detects that the fire drill service group is different from the application service group, it displays the Recreate Fire Drill Service Group panel. Clear the option to recreate the fire drill service group and click **Next**.
- 8 If the wizard detects that the fire drill service group is still online, the Fire Drill Restoration panel is displayed. Review the requirements for restoration and click **Next**.
- 9 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next**.
- 10 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**, and click **Yes** to confirm the deletion.
- 11 The Fire Drill Deletion panel shows the progress of the deletion. Wait until all tasks are complete and then click **Next**.
If errors occur while deleting the fire drill configuration, the wizard will list any incomplete steps so that you can complete them manually.
- 12 The Summary panel is displayed. Click **Finish**.

Microsoft Clustering Solutions

This section contains the following chapters:

- [Microsoft clustering solutions overview](#)
- [Deploying SFW with MSCS](#)
- [Deploying SFW with Microsoft failover clustering](#)
- [Deploying SFW with MSCS in a campus cluster](#)
- [Deploying SFW with Microsoft failover clustering in a campus cluster](#)
- [Deploying SFW and VVR with MSCS](#)
- [Deploying SFW and VVR with Microsoft failover clustering](#)

Microsoft clustering solutions overview

[About Microsoft clustering with high availability](#)

[About Microsoft clustering with Veritas Volume Replicator](#)

[About Microsoft clustering with campus clustering](#)

About Microsoft clustering with high availability

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

A high availability solution maintains continued functioning of applications in the event of computer failure, where data and applications are available using redundant software and hardware. “High availability” can refer to any software or hardware that provides fault tolerance, but generally it has become associated with clustering.

A cluster is a group of independent computers working together as a single system to ensure that mission-critical applications and resources are highly available. The cluster is managed as a single system, shares a common namespace, and is specifically designed to tolerate component failures and to support the addition or removal of components in a way that is transparent to users.

Keeping data and applications functioning 24 hours a day and seven days a week is a requirement for critical applications today. Clustered systems have several advantages, including fault tolerance, high availability, scalability, simplified management, and support for rolling upgrades.

About Microsoft clustering with Veritas Volume Replicator

Microsoft clustering may be used with Veritas Storage Foundation for Windows and Veritas Volume Replicator to provide replication support for your application. Using VVR with Microsoft clustering provides a replicated backup of your application 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.

About Microsoft clustering with campus clustering

Campus clusters are multiple-node clusters that provide protection against disasters. These clusters are in separate buildings (or sites) with mirrored SAN-attached storage located in each building. Typical campus clusters involve two sites; you can use more than two sites for additional redundancy.

In a typical configuration, each node has its own storage array and contains mirrored data of the storage on the other array.

This environment also provides a simpler solution for disaster recovery than a more elaborate Symantec disaster recovery environment with replication software; however, a campus cluster generally stretches a shorter distance than a replication-based solution depending on the hardware.

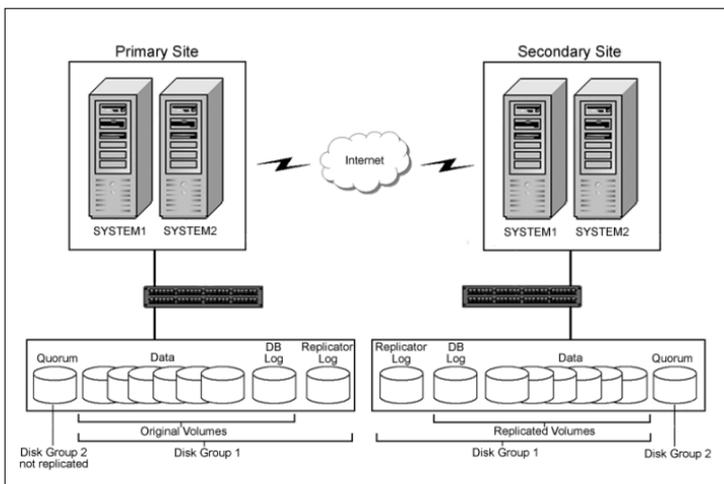
About the SFW-Microsoft clustering-VVR configuration

A typical disaster recovery configuration requires that you have a source host on the primary site and a destination host on the secondary site. The application data is stored on the primary site and replicated to the secondary site by using a tool such as the Veritas Volume Replicator. The primary site provides data and services during normal operation. If a disaster occurs on the primary site and its data is destroyed, a secondary host can take over the role of the primary host to make the data accessible. The application can be restarted on that host.

This Disaster Recovery section includes a SFW-Microsoft clustering-VVR configuration. The configuration is described with a generic database application that includes both data and a database log.

The illustration below shows the SFW HA-VVR configuration with Microsoft clustering. For a SFW-Microsoft clustering-VVR configuration, at least two disk groups are necessary—one for the application and one for the quorum resource volume, which has to be in a separate disk group, as shown in the illustration that follows.

Figure 16-1 SFW-Microsoft clustering-VVR configuration



The quorum volume is not replicated from the primary site to the secondary site. Each site has its own quorum volume. A two-way or four-way mirror is recommended for the quorum volume for redundancy.

Configuring the quorum device for high availability

Either a single basic disk used as a physical disk resource or a volume located on a three-disk SFW cluster disk group can serve as the Microsoft clustering quorum device.

In general, a disk group containing a dedicated, three-way mirrored volume makes an ideal quorum device.

In Microsoft clustering environments, the proper configuration of a quorum device is critical to providing the highest availability with SFW storage.

Using a single disk as the quorum device introduces a nonredundant component into an otherwise highly available system. A failure-tolerant volume used as a quorum device provides a level of availability that is consistent with that of the rest of the cluster.

An SFW cluster disk group containing a volume used as a quorum device should contain that volume only. Any other volumes in that disk group fail over whenever the quorum device changes ownership.

A disk group containing only a three-way mirrored volume makes an ideal quorum device. Such a device tolerates both disk failures, because it is mirrored, and server and interconnect failures, because SFW can import it when the disks and at least one server are running.

For a server to take ownership of a disk group containing the cluster quorum device, SFW must successfully import the disk group, and obtain SCSI reservations on more than half of its disks. Disk groups containing odd numbers of disks are best for use as quorum devices because of this behavior.

Deploying SFW with MSCS

This chapter contains the following topics:

- [“Tasks for deploying SFW with MSCS \(Windows Server 2003\)”](#) on page 609
- [“Reviewing the requirements”](#) on page 612
- [“Reviewing the configuration”](#) on page 615
- [“Configuring the network and storage”](#) on page 616
- [“Establishing an MSCS cluster”](#) on page 618
- [“Installing SFW”](#) on page 618
- [“Creating SFW disk groups and volumes”](#) on page 626
- [“Setting up a group for the application in MSCS”](#) on page 633
- [“Installing the application on cluster nodes”](#) on page 635
- [“Completing the setup of the application group in MSCS”](#) on page 636
- [“Implementing a dynamic quorum resource”](#) on page 637
- [“Verifying the cluster configuration”](#) on page 640

Tasks for deploying SFW with MSCS (Windows Server 2003)

This chapter describes how to install and configure Storage Foundation for Windows with MSCS on Windows Server 2003 in a new installation, using an example two-node active/passive cluster configuration.

For information on using Storage Foundation for Windows with Microsoft failover clustering on Windows Server 2008, see:

[Chapter 18, “Deploying SFW with Microsoft failover clustering”](#) on page 643

The example describes a generic database application in order to present general recommendations that apply to multiple applications. For specific examples of a SFW-MSCS clustering solution see the following:

- *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange*
- *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*

Table 17-1 outlines the high-level objectives for implementing the configuration and the tasks for each objective.

Table 17-1 Task list for deploying SFW with MSCS

Objectives	Tasks
“Reviewing the requirements” on page 612	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites. ■ Review the configuration requirements.
“Configuring the network and storage” on page 616	<ul style="list-style-type: none"> ■ Install the operating system on both nodes. ■ Make necessary networking settings on both nodes.
“Establishing an MSCS cluster” on page 618	<ul style="list-style-type: none"> ■ Refer to Microsoft documentation for instructions on establishing a cluster under MSCS.
“Installing SFW” on page 618	<ul style="list-style-type: none"> ■ Verify the driver signing options for Windows 2003 systems ■ Install SFW ■ Install Cluster Option for Microsoft Cluster Service (MSCS) ■ Restore driver signing options for the Windows 2003 systems
“Creating SFW disk groups and volumes” on page 626	<ul style="list-style-type: none"> ■ In SFW on Node A, create at least two dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum. ■ The disk group for the quorum can be created later, if desired.

Table 17-1 Task list for deploying SFW with MSCS (Continued)

Objectives	Tasks
“Setting up a group for the application in MSCS” on page 633	<ul style="list-style-type: none">■ Create a group within MSCS for the application.■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application on cluster nodes” on page 635	<ul style="list-style-type: none">■ Install the application program files on the local drive of the first node.■ Install files relating to the data and logs on the shared storage.■ Use Move Group to move the cluster resources to the second node.■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node.■ Install the application on the second node.
“Completing the setup of the application group in MSCS” on page 636	<ul style="list-style-type: none">■ Refer to the application documentation for help on creating its resource.■ Establish the appropriate dependencies.■ Test the application group by using the Move Group command to move the cluster resources to the other node.
“Implementing a dynamic quorum resource” on page 637	<ul style="list-style-type: none">■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.■ Make that disk group a Volume Manager Disk Group type resource in the default Cluster Group.■ Change the quorum resource to the dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 640	<ul style="list-style-type: none">■ Use the Move Group command to move the cluster resources to the second node. Move them back to the first node.■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.

Reviewing the requirements

Verify that the following requirements for your configuration are met before starting the Veritas Storage Foundation for Windows installation:

- “Supported software” on page 612
- “Disk space requirements” on page 614
- “System requirements” on page 614

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 SFW with 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

Your system must run one of the following operating systems to install the SFW server software.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 17-5 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required

Table 17-5 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

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 17-6 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

- The configuration described requires shared disks to support applications that migrate between nodes in the 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 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.
- Each system requires 1 GB of RAM for SFW.
- 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.

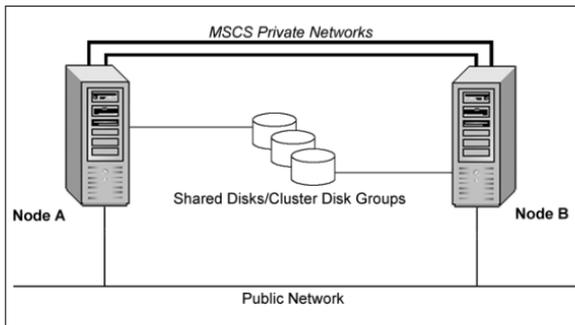
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported. Six network interface cards, three for each server (two each for the private network and one for the public network) are required. You also need a static IP address for the cluster 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.

Reviewing the configuration

The example of a new installation with two servers and one storage array in an active/passive configuration is a typical configuration for a cluster. In an active/passive configuration the active node of the cluster hosts the virtual server and the second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. The example describes a generic database application.

Figure 17-1 Storage Foundation configuration with MSCS and two servers



This configuration does not include DMP. For information about DMP and clustering, see [Chapter 6, "Adding DMP to a clustering configuration"](#) on page 67.

Key points about the configuration:

- An MSCS cluster must be running to install SFW. Therefore, you need to set up the hardware and install the operating system and MSCS on both systems and establish the MSCS 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 active 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*.

- 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 single 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 a quorum disk fails, a mirror on another disk (another plex) takes over and the resource remains online. For 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 your environment is configured to move the quorum disk to a dynamic mirrored quorum volume; this enables you to verify that the application is working in the cluster before adding the dynamic quorum volume.

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 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.

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 619.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 621.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 624.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
See “[Changing the driver signing options](#)” on page 619.
- Moving the Online Groups
See “[Moving the online groups](#)” on page 620.

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 17-7 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 17-7 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 624.
- Completing the SFW Installation
See “[Completing the SFW installation](#)” on page 624.
- Resetting the driver signing options
See “[Resetting the driver signing options](#)” on page 625.

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 619.

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 SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different volume layouts. Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic cluster disk groups](#)
- [Creating dynamic volumes](#)

Planning disk groups and volumes

Decide how you want to organize the disk groups and the number and type of volumes you want to create.

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 which depend on the traffic load.
- The number of disk groups that are needed

The number of disk groups depends on your application and the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application, such as the storage group in Microsoft Exchange, is contained in a single disk group. You will also need a disk group with three disks for the mirrored quorum resource. If possible, use small disks.

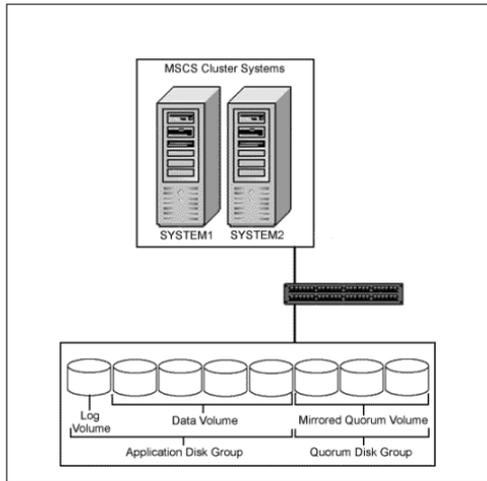
Microsoft recommends 500 MB for the quorum disk.

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.

The following illustration shows a typical setup of shared storage disks for a clustered database application and a dynamic mirrored quorum resource. The log volume is on a separate disk. The log and data volumes are in the application dynamic cluster disk group. The dynamic mirrored quorum is in a separate disk group and has a minimum of two disks, but three are recommended for added fault tolerance.

Figure 17-2 MSCS clustered database with disks for data, the log, and the quorum resource



Creating dynamic cluster disk groups

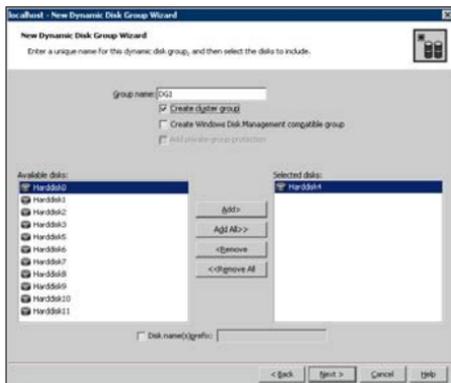
Follow the steps in this section to create one or more disk groups for your application.

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 disk group name (for example, DG1).
- 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 dynamic volumes

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

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.
- 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.

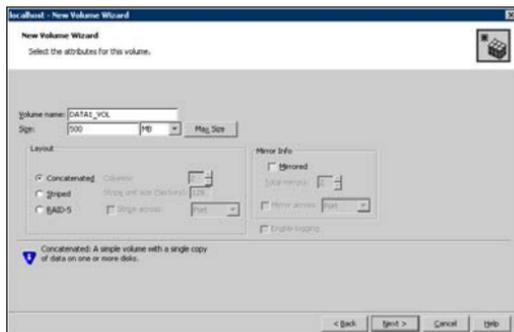


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

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.

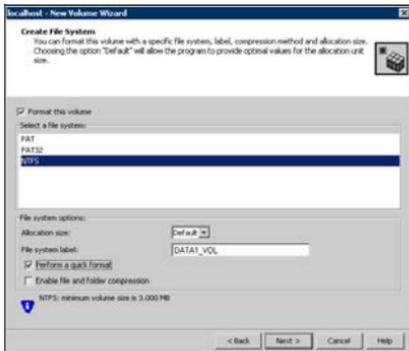


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** 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 additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Setting up a group for the application in MSCS

The next task is to use Cluster Administrator in MSCS to set up a group for the application that will contain the SFW disk group or groups that were created for the application. The SFW disk groups will be added to the MSCS application group as Volume Manager Disk Group resources.

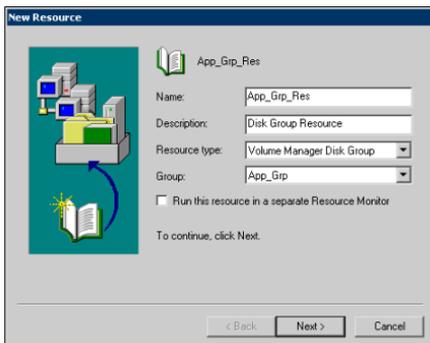
After the application is installed on both nodes and its accompanying files are placed on the shared storage, complete the setup of the application group by adding the application itself as a resource and any other resources that are required. Dependencies need to be set between the resources in the group. Information on this task is included in [“Completing the setup of the application group in MSCS”](#) on page 636.

To set up the application group

- 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**, click **New**, and click **Group**.
- 3 When the New Group dialog box appears, specify a name for the group (for example, App_Grp).
- 4 Click **Next** to continue.
- 5 When the Preferred Owners dialog box appears, make sure that all the preferred owners are added to the **Preferred Owners** list.
- 6 Click **Finish** to create the group.

To create a disk group resource

- 1 From the Cluster Administrator configuration tree, right-click on the MSCS group that you have created (App_Grp) and click **New > Resource**.
- 2 In the New Resource dialog box, specify a name for the disk group resource and a description for the resource, if necessary.



- 3 Select **Volume Manager Disk Group** from the **Resource type** list.
- 4 Select **App_Grp** from the **Group** list.
- 5 Click **Next**.
- 6 In the Possible Owners dialog box, click **Next**.
- 7 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a Disk Group resource.
- 8 When the Volume Manager Disk Group Parameters dialog box appears, select the disk group.
- 9 Click **Finish**.
- 10 Click **OK**.
- 11 Bring the resources online.

Installing the application on cluster nodes

The application program files need to be installed on the same local drive on all the cluster nodes. The application data and log files or other files related to the application data are installed on the shared storage.

Pointers 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 SQL Server, install on both nodes at once.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Do not to accept the default locations for the application data and log files. Instead, set the paths for these files to the drive letters or mount points of the volumes created in [“Creating dynamic volumes”](#) on page 630.

Pointers for installing the application on the second node

- Use the **Move Group** command to move the cluster resources to the second node.
- Verify that the volumes on shared storage can be accessed from the second node using the same drive letters or mount points that were assigned when they were created on the first node. To change a drive letter or mount point, see [“To add or change a drive letter or mount point”](#) on page 635.
- 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. You 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**. The Drive Letter and Paths window appears.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.

- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Change the drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

Note: A mount point is also referred to as a "drive path."

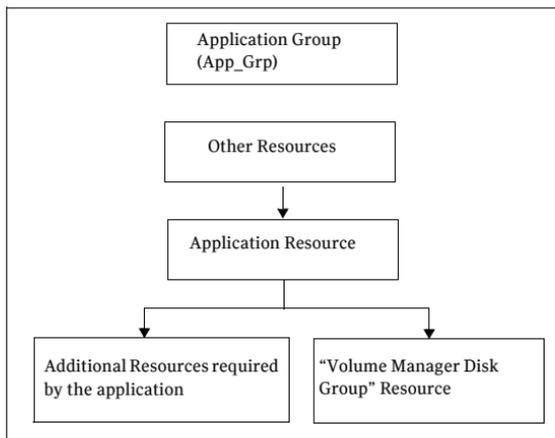
- 6 To change a mount point, you must remove it and recreate it (step 5). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

Completing the setup of the application group in MSCS

The additional steps in this section make the application group functional in MSCS. The application resource needs to be added, as well as any other resources that are associated with the application. Also, dependencies need established for the resources. This section presents a high-level summary of the process for completing the application group setup.

- Before creating the application resource, make sure that all the other resources, including the disk group resource and any additional application resources, are online.
- Refer to the application documentation for information on creating its resource and additional resources that may be required. You will need to create an IP address resource and a Network name resource in addition to the Volume Manager Disk Group resource you created earlier. Note that when creating the application resource, on the Dependencies screen, select the **Volume Manager Disk Group** resource from "Available Resources" and add it to "Resource Dependencies."
- The following dependency chart indicates the dependencies that are established.

Figure 17-3 Application group dependencies



- **Testing:** After the application group is set up, test it by using the **Move Group** command to move the cluster resources to another node and then move them back.

Implementing a dynamic quorum resource

Although Symantec recommends implementing a dynamic quorum resource in order to take full advantage of the Storage Foundation functionality, it is not a required task:

- [“Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume”](#) on page 638
- [“Creating the quorum resource for the cluster group”](#) on page 638
- [“Changing the quorum resource to a dynamic mirrored quorum resource”](#) on page 640

Note: If you are using DMP, you must create a dynamic quorum resource in order for the groups to fail over properly.

Note: If you are planning to create a dynamic mirrored quorum for your Volume Manager Disk (VMDG) resource, then make your ClusSvc resource dependent on the VxBridge service. Refer to Microsoft documentation for details on creating resource dependencies. This dependency is required only in cases when you will reboot all the cluster nodes at the same time

Creating a dynamic cluster disk group for the Quorum Resource with mirrored volume

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** checkbox, and specify three mirrors. For full details about creating cluster disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 626.

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.

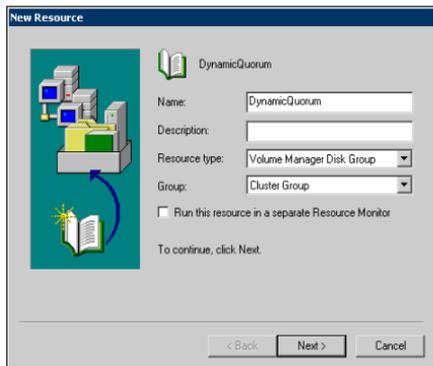
Creating the quorum resource for the cluster group

Create a quorum resource for the cluster group to protect the quorum during failover.

To create the quorum resource

- 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 Create the quorum resource, right-click the **Cluster Group**, click **New**, and click **Resource**.



- 3 When the New Resource dialog box appears, specify a name for the quorum resource (QuorumDG) and, if necessary, add a description about the resource.
- 4 Select **Volume Manager Disk Group** from the **Resource type** list.
- 5 Select **Cluster Group** from the **Group** list.
- 6 Click **Next**.
- 7 In the Possible Owners dialog box, click **Next**.
- 8 In the Dependencies dialog box, click **Next**. You do not need to set dependencies for a quorum resource.
- 9 When the Volume Manager Disk Group Parameters dialog box appears, select the disk group.
- 10 Click **Finish**.
- 11 Bring the newly added resource online.

Changing the quorum resource to a dynamic mirrored quorum resource

The last step in this process is to change the quorum resource to a dynamic mirrored quorum.

To change the quorum resource to a dynamic mirrored quorum

- 1 From Cluster Administrator, right-click the cluster name in the configuration tree, and click **Properties**.
- 2 Select the Quorum tab of the Properties window.
- 3 Select the name of the dynamic quorum disk group resource added in ["Creating the quorum resource for the cluster group"](#) on page 638.
- 4 Click **OK**.

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 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

This chapter covers the following topics:

- [“Tasks for deploying SFW with Microsoft failover clustering \(Windows Server 2008\)”](#) on page 644
- [“Reviewing the requirements”](#) on page 646
- [“Reviewing the configuration”](#) on page 649
- [“Configuring the storage hardware and network”](#) on page 650
- [“Establishing a failover cluster”](#) on page 652
- [“Installing SFW”](#) on page 653
- [“Creating SFW disk groups and volumes”](#) on page 659
- [“Creating a group for the application in the failover cluster”](#) on page 667
- [“Installing the application on cluster nodes”](#) on page 668
- [“Completing the setup of the application group in the failover cluster”](#) on page 670
- [“Implementing a dynamic quorum resource”](#) on page 671
- [“Verifying the cluster configuration”](#) on page 673

Tasks for deploying SFW with Microsoft failover clustering (Windows Server 2008)

This chapter describes how to install and configure Storage Foundation for Windows with Microsoft failover clustering (Windows Server 2008) in a new installation, using a two-node active/passive cluster configuration as an example.

For information on using Storage Foundation for Windows with MSCS on Windows Server 2003, see:

[Chapter 17, “Deploying SFW with MSCS” on page 609](#)

The example describes a generic database application in order to present general recommendations that apply to multiple applications. For specific examples of an SFW-Microsoft failover cluster solution, see the following:

- *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft Exchange*
- *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*

[Table](#) outlines the high-level objectives for implementing the configuration and the tasks for each objective:

Task list for deploying SFW with Microsoft failover clustering

Objectives	Tasks
“Reviewing the requirements” on page 646	<ul style="list-style-type: none">■ Verify hardware and software prerequisites.■ Review the configuration requirements.
“Reviewing the configuration” on page 649	<ul style="list-style-type: none">■ Understanding the configuration for the failover cluster
“Configuring the storage hardware and network” on page 650	<ul style="list-style-type: none">■ Install the operating system on both nodes.■ Make necessary networking settings on both nodes.
“Establishing a failover cluster” on page 652	<ul style="list-style-type: none">■ Refer to Microsoft documentation for instructions on establishing a cluster under Microsoft failover clustering.
“Installing SFW” on page 653	<ul style="list-style-type: none">■ Install SFW■ Install Cluster Option for Microsoft Cluster Service (MSCS)

Task list for deploying SFW with Microsoft failover clustering (Continued)

Objectives	Tasks
“Creating SFW disk groups and volumes” on page 659	<ul style="list-style-type: none">■ In SFW on Node A, create at least two dynamic cluster disk groups on the storage—one or more for the application data files and one for the mirrored quorum.■ The disk group for the quorum can be created later, if desired.
“Creating a group for the application in the failover cluster” on page 667	<ul style="list-style-type: none">■ Create a group within the failover cluster for the application.■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application on cluster nodes” on page 668	<ul style="list-style-type: none">■ Install the application program files on the local drive of the first node.■ Install files relating to the data and logs on the shared storage.■ Move the cluster resources to the second node.■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node.■ Install the application on the second node.
“Completing the setup of the application group in the failover cluster” on page 670	<ul style="list-style-type: none">■ Refer to the application documentation for help on creating its resource.■ Establish the appropriate dependencies.■ Test the application group by moving the cluster resources to the other node.
“Implementing a dynamic quorum resource” on page 671	<ul style="list-style-type: none">■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.■ Make that disk group a Volume Manager Disk Group type resource in the default Cluster Group.■ Configure the quorum resource.
“Verifying the cluster configuration” on page 673	<ul style="list-style-type: none">■ Move the cluster resources to the second node. Move them back to the first node.■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.

Reviewing the requirements

Verify that the following requirements for your configuration are met before starting the Storage Foundation for Windows installation:

- “Supported software for Microsoft failover clustering with SFW” on page 646
- “Disk space requirements” on page 647
- “System requirements” on page 647

Supported software for Microsoft failover clustering with SFW

The following software is supported for deploying SFW and Microsoft clustering on Windows Server 2008:

The following table shows supported Windows Server 2008 operating systems.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

Table 18-1 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition	

Table 18-1 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008 R2 for IA systems	IA64		
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 18-2 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

- The configuration described requires shared disks to support applications that migrate between nodes in the cluster.
- SCSI or Fibre Channel host bus adapters (HBAs) can be used to access shared storage.
- Microsoft 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 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.
- Each system requires 1 GB of RAM for SFW.
- Systems to be clustered must be configured as part of a Windows Server 2008 domain. Each system in an Microsoft failover cluster must be in the same domain and must be using the same operating system version.

Reviewing the requirements

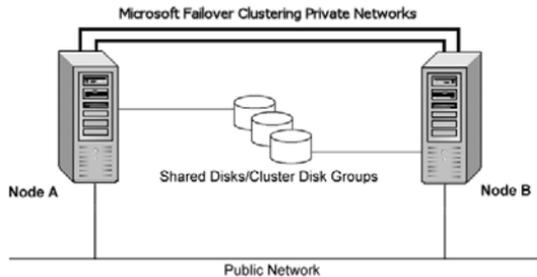
- To install and work with the SFW and Microsoft clustering software, you must have an account with Administrator privileges. You must also have a license key to install SFW.
- Using static IP addresses for the public network and private network cards is highly recommended. DHCP is not supported for VVR. Six network interface cards, three for each server (two each for the private network and one for the public network) are required. 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.

Reviewing the configuration

The example of a new installation with two servers and one storage array in an active/passive configuration is a typical configuration for a cluster. In an active/passive configuration the active node of the cluster hosts the virtual server and the second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. The example describes a generic database application.

Figure 18-1 Storage Foundation configuration with Microsoft failover clustering and two servers



This configuration does not include DMP. For information about DMP and clustering, see [“Adding DMP to a clustering configuration”](#) on page 67.

Key points about the configuration:

- A Microsoft failover cluster must be running before you install SFW. Therefore, you need to set up the hardware and install the operating system on both 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 active cluster resources to the other node, and install on the now inactive node.
- 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. For 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.

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 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

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 654.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 654.

- Performing post-installation tasks
See ["Post-installation tasks"](#) on page 658.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups
See ["Moving the online groups"](#) on page 654.

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 658.
- Completing the SFW installation
See “[Completing the SFW installation](#)” on page 659.

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 653.

Creating SFW disk groups and volumes

Use Veritas Storage Foundation for Windows to create disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different volume layouts. Configuring disk groups and volumes involves the following tasks:

- [Planning disk groups and volumes](#)
- [Creating dynamic cluster disk groups](#)
- [Creating dynamic volumes](#)

Planning disk groups and volumes

Decide how you want to organize the disk groups and the number and type of volumes you want to create.

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 which depend on the traffic load.
- The number of disk groups that are needed

The number of disk groups depends on your application and the planned organization of the data. The application program files need to be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application, such as the storage group in Microsoft Exchange, is contained in a single disk group.

You will also need to create a disk group with three disks and a mirrored volume for the quorum resource. If possible, use small disks. Microsoft recommends 500 MB for the quorum disk. See “[Implementing a dynamic quorum resource](#)” on page 671.

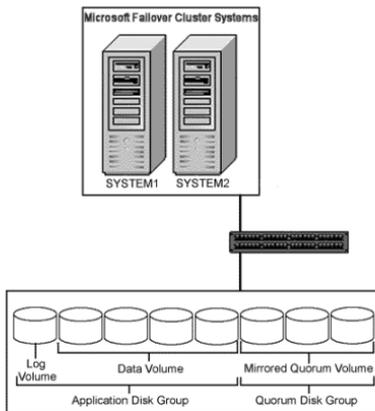
Recommendations:

- Use mirrored volumes for logs.

- Use striped or mirrored striped volumes for data.

The following illustration shows a typical setup of shared storage disks for a clustered database application and a dynamic mirrored quorum resource. The log volume is on a separate disk. The log and data volumes are in the application cluster disk group. The dynamic mirrored quorum is in a separate disk group and has a minimum of two disks, but three are recommended for added fault tolerance.

Figure 18-2 Microsoft failover clustered database with disks for data, the log, and the quorum resource



Creating dynamic cluster disk groups

Follow the steps in this section to create one or more disk groups for your application.

Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

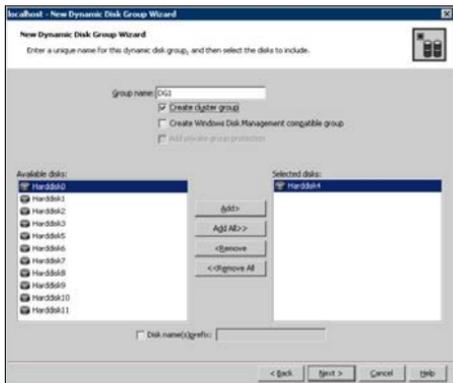
Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

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 disk group name (for example, DG1).
- 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 dynamic volumes

Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

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.
- 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.

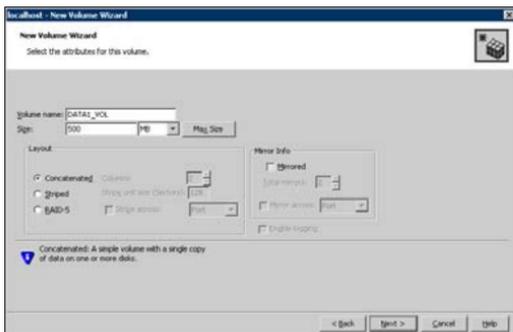


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

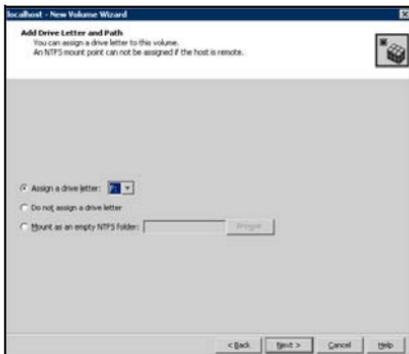
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the volume attributes.

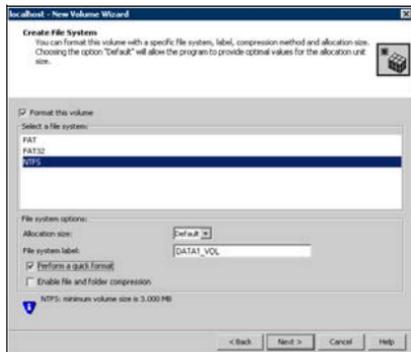


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** 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 additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Creating a group for the application in the failover cluster

After you create SFW disk groups and volumes for the application, use the Failover Cluster Management tool to set up a cluster group for the application. You then add Volume Manager Disk Group resources for the SFW disk groups that you created for the application.

After the application is installed on both nodes and its accompanying files are placed on the shared storage, you will do additional steps to complete the setup of the application group.

See “[Completing the setup of the application group in the failover cluster](#)” on page 670.

To set up the application cluster group

- 1 Launch Failover Cluster Management by selecting **Start > Administrative Tools > Failover Cluster Management**. Connect to the appropriate cluster through the console.
- 2 Create a new group by selecting the **Services and Applications** node from the tree that is displayed in the left hand pane. Right-click and select **More Actions > Create Empty Service or Application**. An empty group named New service or application is created.
- 3 Specify a name for the group by right-clicking it and selecting **Rename** from the drop down menu .
- 4 Type the name of the new group (for example, App_Grp) in the **Name** field. You can now add Volume Manager Disk Group resources to the application group.

To create a Volume Manager Disk Group resource for the application

- 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 In the left pane of Failover Cluster Management, right-click the application cluster group (for example, App_Grp) and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 3 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** to open its Properties dialog box.
- 4 On the General tab of the Properties dialog box, type a name for the resource.
For example, type APP_DG_RES.

- 5 On the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the application (for example, DG1), and click **OK** to close the dialog box.
- 6 Right-click the newly named resource and select **Bring this resource online**.
- 7 If you created more than one disk group for the application, repeat this procedure to add another Volume Manager Disk Group resource for another disk group.

Installing the application on cluster nodes

Install the application program files on the same local drive on all the cluster nodes. Install the application data and log files or other files related to the application data on the shared storage.

Pointers 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 install on both nodes at once.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Do not accept the default locations for the application data and log files. Instead, set the paths for these files to the drive letters or mount points of the volumes created in [“Creating dynamic volumes”](#) on page 663.

Pointers for installing the application on the second node

- In Failover Cluster Management, move the cluster resources to the second node.
- Verify that the volumes on shared storage can be accessed from the second node using the same drive letters or mount points that were assigned when they were created on 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. You 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**. The Drive Letter and Paths window appears. You can choose from the following:
 - To add a drive letter, click **Add**. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
 - To change a drive letter, click **Modify**. The **Assign a drive letter** drop-down list becomes available. Change the drive letter and click **OK**.
 - To add a mount point, click **Add**, 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, and click **OK** to mount the volume.
 - To change a mount point, you must remove it and then select the Add option to add it back. To remove it, select it in the Drive Letter and Paths window and click **Remove**.

Completing the setup of the application group in the failover cluster

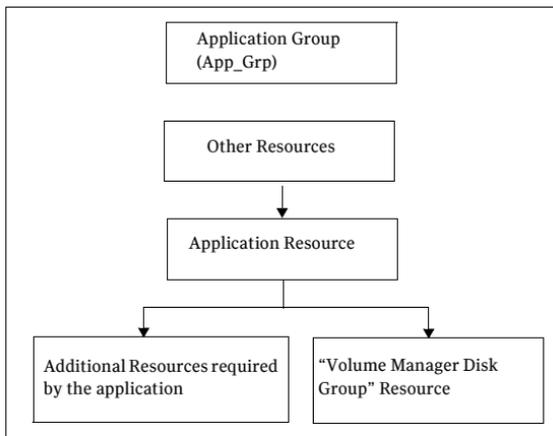
The additional steps in this section make the application group functional in the failover cluster. The application resource needs to be added, as well as any other resources that are associated with the application. Also, dependencies need to be established for the resources. The following list presents a high-level summary of the process for completing the application group setup.

- Before creating the application resource, make sure that all the other application resources, including the Volume Manager Disk Group resource and any additional application resources, are online.
- Refer to the application documentation for information on creating its resource and additional resources that may be required, such as the IP address resource.

When creating the application resource, add the Volume Manager Disk Group resource as a resource dependency.

- The following dependency chart indicates the dependencies that are established.

Figure 18-3 Application group dependencies



- **Testing:** After the application group is set up, test it by moving the cluster resources to another node and then move them back.

Implementing a dynamic quorum resource

Although Symantec recommends implementing a dynamic quorum resource in order to take full advantage of the Storage Foundation functionality, it is not a required task. To implement a dynamic quorum resource, complete the following tasks:

- [“Creating a dynamic cluster disk group and a mirrored volume for the quorum resource”](#) on page 671
- [“Creating the quorum resource for the cluster group”](#) on page 673
- [“Changing the quorum resource to a dynamic mirrored quorum resource”](#) on page 673

Note: If you are using DMP, you must create a dynamic quorum resource in order for the groups to fail over properly.

Note: If you are planning to create a dynamic mirrored quorum for your Volume Manager Disk (VMDG) resource, then make your ClusSvc resource dependent on the VxBridge service. Refer to Microsoft documentation for details on creating resource dependencies. This dependency is required only in cases when you will reboot all the cluster nodes at the same time

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 in the New Volume wizard, select the **Concatenated** layout, select the **Mirrored** checkbox, and specify three mirrors.

For full details about creating cluster disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 659.

Implementing a dynamic quorum resource

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.

Creating the quorum resource for the cluster group

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

Verifying the cluster configuration

You can verify your installation by moving the cluster group between nodes to see if it fails over properly. The ultimate test of the cluster's failover capability involves shutting down the node that is currently online and bringing it back up after the cluster fails over to the other node.

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 with MSCS in a campus cluster

This chapter discusses the following topics:

- [“Tasks for deploying SFW with MSCS in a campus cluster \(Windows Server 2003\)”](#) on page 675
- [“Reviewing the requirements”](#) on page 678
- [“Reviewing the configuration”](#) on page 681
- [“Configuring the network and storage”](#) on page 688
- [“Establishing an MSCS cluster”](#) on page 690
- [“Installing SFW”](#) on page 693
- [“Creating disk groups and volumes”](#) on page 700
- [“Changing the quorum resource to a dynamic quorum resource”](#) on page 708
- [“Setting up a group for the application in MSCS”](#) on page 711
- [“Installing the application on the cluster nodes”](#) on page 711
- [“Completing the setup of the application group in MSCS”](#) on page 713
- [“Verifying the cluster configuration”](#) on page 714

Tasks for deploying SFW with MSCS in a campus cluster (Windows Server 2003)

This chapter presents a Microsoft clustering (MSCS) campus clustering example using a two-node cluster. The operating system is Windows Server 2003.

For information on deploying SWF with Microsoft failover clustering, which runs under Windows Server 2008, see:

[Chapter 20, “Deploying SFW with Microsoft failover clustering in a campus cluster” on page 717.](#)

[Table 19-1](#) lists the high-level objectives for deploying SFW with MSCS in a campus cluster, as well as the tasks within each objective:

Table 19-1 Task list for deploying SFW with MSCS in a campus cluster

Objectives	Tasks
“Reviewing the requirements” on page 678	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites.
“Reviewing the configuration” on page 681	<ul style="list-style-type: none"> ■ Review the configuration requirements. ■ Overview of MSCS campus cluster, and recovery scenarios
“Configuring the network and storage” on page 688	<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.
“Establishing an MSCS cluster” on page 690	<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.
“Installing SFW” on page 693	<ul style="list-style-type: none"> ■ Install SFW on Node A (Node B active). ■ Install SFW on Node B (Node A active).

Table 19-1 Task list for deploying SFW with MSCS in a campus cluster

Objectives	Tasks
“Creating disk groups and volumes” on page 700	<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.
“Changing the quorum resource to a dynamic quorum resource” on page 708	<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.
“Setting up a group for the application in MSCS” on page 711	<ul style="list-style-type: none"> ■ Create a group within MSCS for the application. ■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application on the cluster nodes” on page 711	<ul style="list-style-type: none"> ■ Install the application program files on the local drive of the first node. ■ Install files relating to the data and logs on the shared storage. ■ Move the cluster resources to the second node. ■ Make sure that the volumes on the second node have the same drive letters or mount points as they had on the first node. ■ Install the application on the second node.
“Completing the setup of the application group in MSCS” on page 713	<ul style="list-style-type: none"> ■ Refer to the application documentation for help on creating its resource. ■ Establish the appropriate dependencies. ■ Test the application group by moving the cluster resources to the other node.
“Verifying the cluster configuration” on page 714	<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 SFW with 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

Your system must run one of the following operating systems to install the SFW server software.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 19-5 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required

Table 19-5 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

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) can be used 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.
- 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 19-6 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.

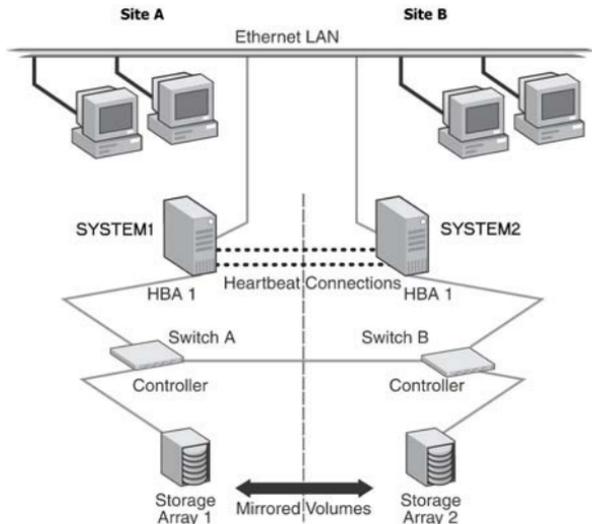
Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with MSCS or for recovery scenarios, see

- [“Overview of campus clustering with MSCS”](#) on page 682
- [“MSCS campus cluster failure scenarios”](#) on page 683

Figure 19-1 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 the DMP chapter, “[Adding DMP to a clustering configuration](#)” on page 67.

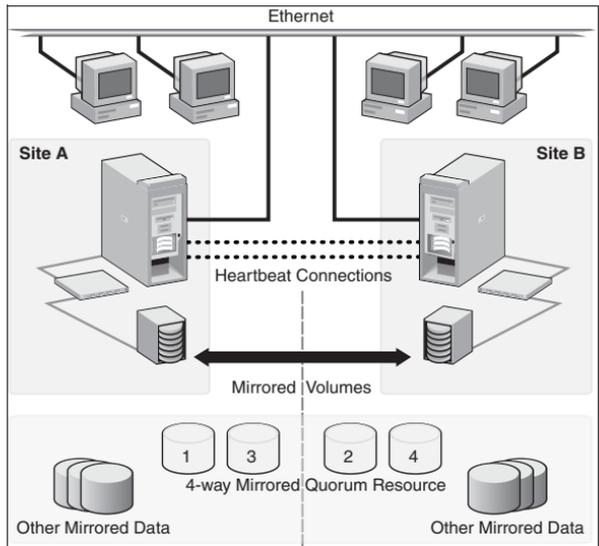
When you are installing SFW and MSCS together, remember the following:

- 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.
- 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 19-2](#) on page 683 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 19-2 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 687.

Table 19-7 lists failure situations and the outcomes that occur:

Table 19-7 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. There will be a temporary service interruption 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 will simply route the heartbeat packets through the public network.

Table 19-7 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
Public IP Network Failure May mean that the public NIC or LAN network has failed.	Failover. Mirroring continues.	When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.
Public and Private IP or Network Failure May mean that the LAN network, including both private and public NIC connections, has failed.	No interruption of service. No Public LAN access. Mirroring continues.	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.
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.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	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.
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.

Table 19-7 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
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 690
- [“Configuring the shared storage and creating a partition for the Cluster quorum disk”](#) on page 691
- [“Creating the first node of the cluster on Server A”](#) on page 691
- [“Installing and configuring the operating system and MSCS on Server B”](#) on page 691
- [“Connecting the two nodes”](#) on page 691
- [“Creating the second node of the cluster on Server B”](#) on page 692
- [“Verifying the cluster configuration”](#) on page 692.

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 690.

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.

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 693.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 695.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 698.

Pre-installation tasks

Perform the following pre-installation tasks:

- Changing the driver signing options
See “[Changing the driver signing options](#)” on page 693.
- Moving the Online Groups
See “[Moving the online groups](#)” on page 695.

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 19-8 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 698.
- Completing the SFW Installation
See “[Completing the SFW installation](#)” on page 699.
- Resetting the driver signing options
See “[Resetting the driver signing options](#)” on page 699.

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 693.

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

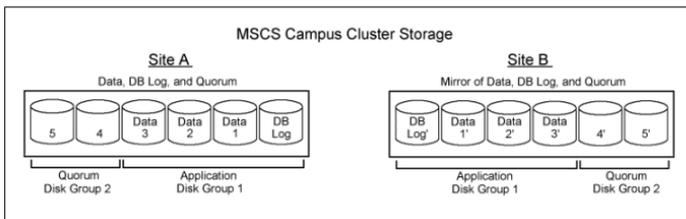
Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

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.

[Figure 19-3](#) 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 19-3](#), 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 19-3 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 702
- [“Creating a volume”](#) on page 704

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

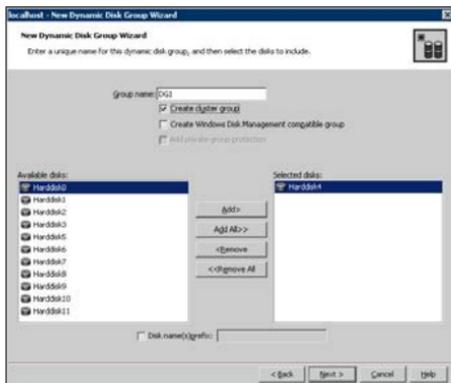
Use the following procedure to create a dynamic cluster disk group.

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 disk group name (for example, DG1).
- 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.

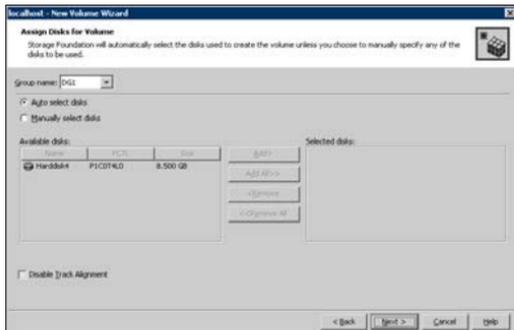
Proceed to create the appropriate volumes on each disk.

Creating a volume

Use the following procedure to create dynamic volumes.

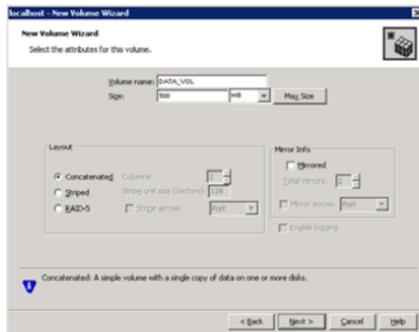
To create dynamic volumes

- 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.
- 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.



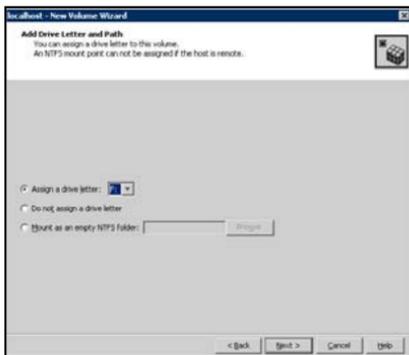
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:

- 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.



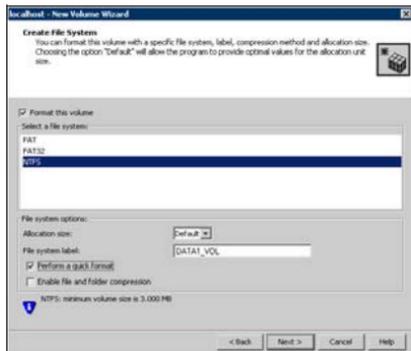
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.
- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that 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**.

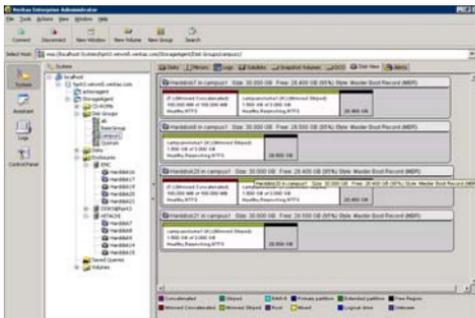
12 Create an NTFS file system.



- 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 additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Figure 19-4 View of disks with volumes in VEA Console



Changing the quorum resource to 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 708
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 709
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 711

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 700.

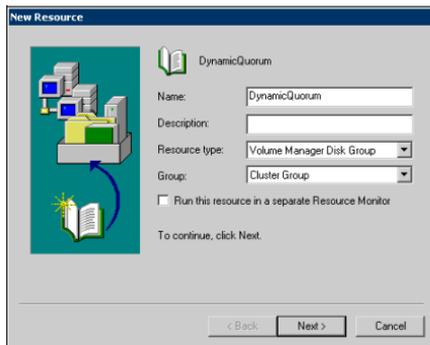
Note: If you add other volumes to this disk group, any failures related to their operation can cause disruptive failovers of the quorum volume. If a volume in the group experiences a high level of read/write activity, failovers may result from delayed access to the quorum volume by MSCS.

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 on that disk group and select **New > Resource**. The New Resource window appears.

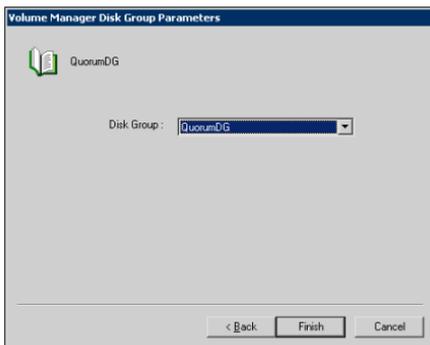


- 3 On the New Resource, window, do the following:
 - Specify a name for the disk group resource in the **Name** field, such as “QuorumDG.”

- 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**.

Setting up a group for the application in MSCS

Using MSCS, you set up a group for the application that contains the SFW disk group or groups that were created for the application.

Add the appropriate SFW disk groups as resources to the application group. You must add the SFW disk groups as the following resource type:

Volume Manager Disk Group

After you install the application on both nodes and place its accompanying files on the shared storage (which in this case is shared across two storage arrays), complete the setup of the application group by adding the application itself as a resource and any other resources that are required. You need to set dependencies between the resources in the group as described in [“Completing the setup of the application group in MSCS”](#) on page 713.

You cannot add the application resource until after the application is installed on both nodes.

Installing the application on the cluster nodes

You must install the application program files on the same local drive of all the cluster nodes. You install the application data and log files (or other files related to the application data) on the shared storage.

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 the application, remember not to accept the default location for the application data and log files. Instead, click to browse to the dynamic volumes that were 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. To change a drive letter or mount point, see [“To add or change a drive letter or mount point”](#) on page 712.
- 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.”

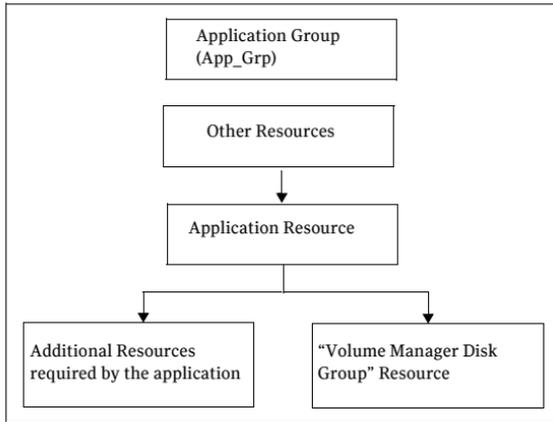
- To change a mount point, you must remove it and add it again. (See the bullet above). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.
- Click **OK**.

Completing the setup of the application group in MSCS

To make the application group functional in MSCS, you need to add the application resource, as well as any other resources that are associated with the application. Also, you need to establish dependencies for the resources. This section summarizes the process for completing the application group setup.

- 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 its resource and additional resources that may be required, see the application documentation. You may need to create an IP address resource and a network name resource in addition to the Volume Manager Disk Group resource that you created earlier.
Make sure that you select the appropriate disk group resource as the storage resource on which the application resource is dependent.
- [Figure 19-5](#) on page 714 shows the application group dependencies.

Figure 19-5 Application group dependencies



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 in a campus cluster

This chapter covers the following topics:

- [“Tasks for deploying SFW with Microsoft failover clustering in a campus cluster \(Windows Server 2008\)”](#) on page 718
- [“Reviewing the requirements”](#) on page 720
- [“Reviewing the configuration”](#) on page 723
- [“Configuring the network and storage”](#) on page 730
- [“Establishing a Microsoft failover cluster”](#) on page 732
- [“Installing SFW”](#) on page 735
- [“Creating disk groups and volumes”](#) on page 740
- [“Implementing a dynamic quorum resource”](#) on page 750
- [“Setting up a group for the application in the failover cluster”](#) on page 752
- [“Installing the application on the cluster nodes”](#) on page 754
- [“Completing the setup of the application group in the cluster”](#) on page 755
- [“Verifying the cluster configuration”](#) on page 757

Tasks for deploying SFW with Microsoft failover clustering in a campus cluster (Windows Server 2008)

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 20-1 Task list for deploying SFW with Microsoft failover clustering in a campus cluster

Objectives	Tasks
“Reviewing the requirements” on page 720	<ul style="list-style-type: none">■ Verify hardware and software prerequisites.
“Reviewing the configuration” on page 723	<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 730	<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 732	<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 735	<ul style="list-style-type: none">■ Install SFW on Node A (Node B active).■ Install SFW on Node B (Node A active).

Table 20-1 Task list for deploying SFW with Microsoft failover clustering in a campus cluster (Continued)

Objectives	Tasks
“Creating disk groups and volumes” on page 740	<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 750	<ul style="list-style-type: none"> ■ If not done earlier, create a dynamic disk group for the quorum with a mirrored volume. ■ Add the volume manager disk group for the quorum. ■ Change the quorum resource to the dynamic mirrored quorum resource.
“Setting up a group for the application in the failover cluster” on page 752	<ul style="list-style-type: none"> ■ Create a group within failover clustering for the application. ■ Include the cluster disk group or groups for the application as Volume Manager. Disk Group type resources in the group.
“Installing the application on the cluster nodes” on page 754	<ul style="list-style-type: none"> ■ Install the application program files on the local drive of the first node. ■ Install files relating to the data and logs on the shared storage. ■ Move the cluster resources to the second node. ■ Make sure that the volumes on the second node have the same drive letters or mount points as they had on the first node. ■ Install the application on the second node.
“Completing the setup of the application group in the cluster” on page 755	<ul style="list-style-type: none"> ■ Refer to the application documentation for help on creating its resource. ■ Establish the appropriate dependencies. ■ Test the application group by moving the cluster resources to the other node.
“Verifying the cluster configuration” on page 757	<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 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 shows supported Windows Server 2008 operating systems.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

Table 20-2 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			

Table 20-2 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition	
Windows Server 2008 R2 for IA systems	IA64		
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.
- 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.

Reviewing the requirements

- 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 20-3 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

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.

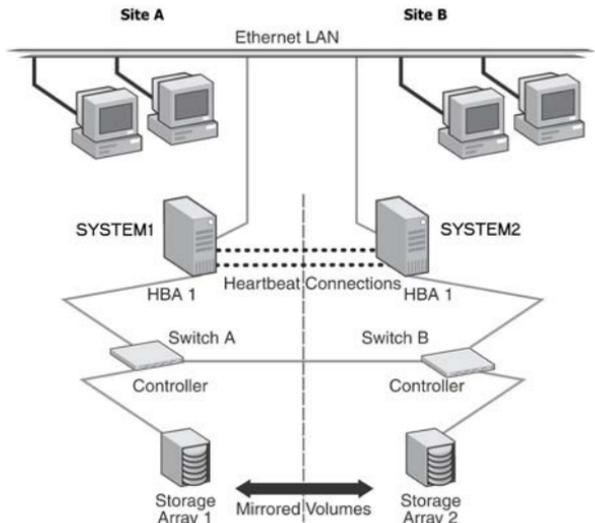
Reviewing the configuration

This configuration example describes a two-node campus cluster with each node at a separate site.

For an overview of campus clusters with Microsoft clustering or for recovery scenarios, see the following:

- “[Overview of campus clustering with Microsoft clustering](#)” on page 724
- “[Campus cluster failure with Microsoft clustering scenarios](#)” on page 725

Figure 20-1 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 the DMP chapter, “[Adding DMP to a clustering configuration](#)” on page 67.

When you are installing SFW and Microsoft clustering together, remember the following:

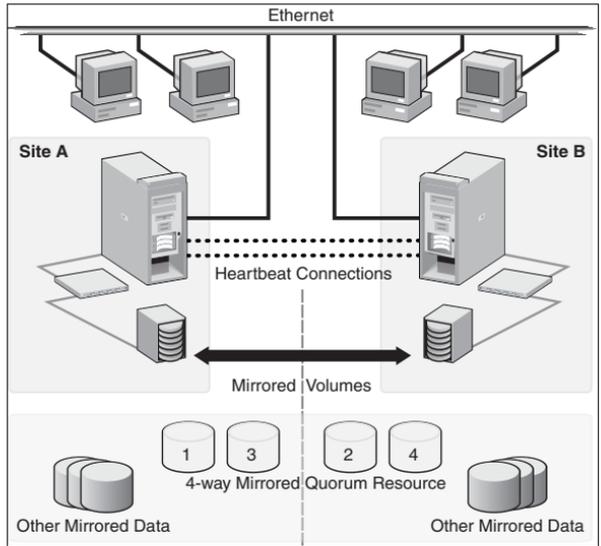
- A cluster using Microsoft clustering must be running to install SFW. You need to set up the hardware and install the operating system and Microsoft clustering on all systems and establish the failover cluster before installing SFW. Installing SFW requires a reboot, but a reboot on the active cluster 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.
- 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 20-2](#) on page 725 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 20-2 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 729.

Table 20-4 lists failure situations and the outcomes that occur:

Table 20-4 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. Service is temporarily interrupted 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.
Private IP Heartbeat Network Failure May mean that the private NICs or the connecting network cables failed.	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 20-4 List of failure situations and possible outcomes (Continued)

Failure Situation	Outcome	Comments
Public IP Network Failure May mean that the public NIC or LAN network has failed.	Failover. Mirroring continues.	When the public NIC on the active node, or public LAN fails, clients cannot access the active node, and failover occurs.
Public and Private IP or Network Failure May mean that the LAN network, including both private and public NIC connections, has failed.	No interruption of service. No Public LAN access. Mirroring continues.	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.
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.	No interruption of service. Disks on the same node are functioning. Mirroring is not working.	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.
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 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.

Table 20-4 List of failure situations and possible outcomes (Continued)

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 onlined 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 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 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.

Establishing a Microsoft failover cluster

- 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.

Establishing a Microsoft failover cluster**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 735.
- Installing the product
See “[Installing Veritas Storage Foundation for Windows](#)” on page 736.
- Performing post-installation tasks
See “[Post-installation tasks](#)” on page 739.

Pre-installation tasks

Perform the following pre-installation tasks:

- Moving the online groups
See “[Moving the online groups](#)” on page 735.

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 740.
- Completing the SFW installation
See “[Completing the SFW installation](#)” on page 740.

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 735.

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

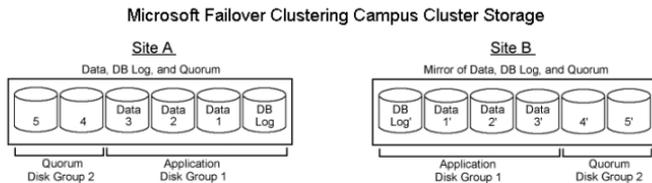
Note: Each volume should be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

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.

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 20-3 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 743
- “Creating a volume” on page 745

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.

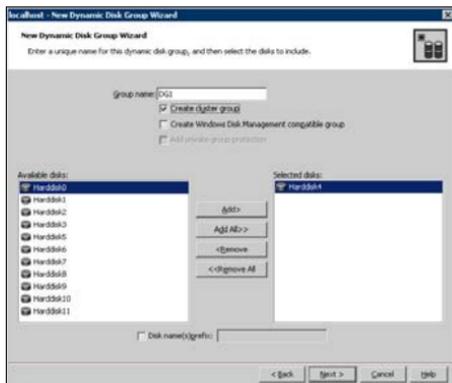
Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

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 disk group name (for example, DG1).
- 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.

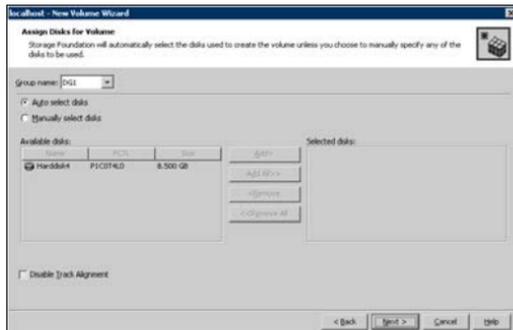
Proceed to create the appropriate volumes on each disk.

Creating a volume

Use the following procedure to create dynamic volumes.

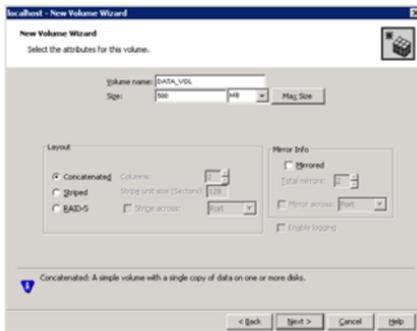
To create dynamic volumes

- 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.
- 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.



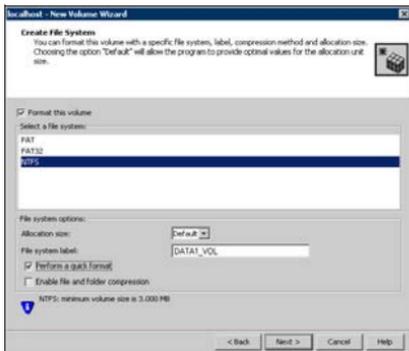
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume.
- If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that 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**.

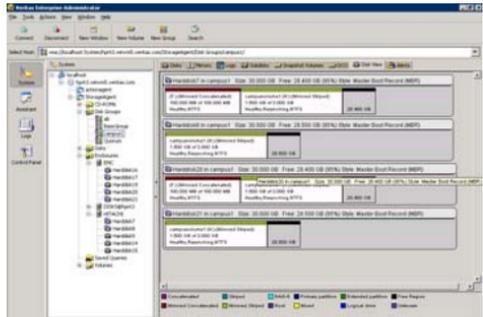
12 Create an NTFS file system.



- 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 additional volumes.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Figure 20-4 View of disks with volumes in VEA Console



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 750
- [“Adding the volume manager disk group for the quorum”](#) on page 750
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 751

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 740.

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 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**.

Setting up a group for the application in the failover cluster

- 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.

Setting up a group for the application in the failover cluster

After you create SFW disk groups and volumes for the application, use the Failover Cluster Management tool to set up a cluster group for the application. You then add Volume Manager Disk Group resources for the SFW disk groups that you created for the application.

After the application is installed on both nodes and its accompanying files are placed on the shared storage, you will do additional steps to complete the setup of the application group.

See “[Completing the setup of the application group in the cluster](#)” on page 755.

To set up the application cluster group

- 1 Launch Failover Cluster Management by selecting **Start > Administrative Tools > Failover Cluster Management**. Connect to the appropriate cluster through the console.
- 2 Create a new group by selecting the **Services and Applications** node from the tree that is displayed in the left hand pane. Right-click and select **More Actions > Create Empty Service or Application**. An empty group named New service or application is created.
- 3 Specify a name for the group by right-clicking it and selecting **Rename** from the drop down menu.
- 4 Type the name of the new group (for example, `App_Grp`) in the **Name** field. You can now add Volume Manager Disk Group resources to the application group.

To create a Volume Manager Disk Group resource for the application

- 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 In the left pane of Failover Cluster Management, right-click the application cluster group (for example, App_Grp) and select **Add a resource > More resources > Add Volume Manager Disk Group**.
- 3 In the center panel under Disk Drives, double-click **New Volume Manager Disk Group** to open its Properties dialog box.
- 4 On the General tab of the Properties dialog box, type a name for the resource.
For example, type APP_DG_RES.
- 5 On the Properties tab, in the Disk Group Name field, type the name of the disk group you previously created for the application (for example, DG1), and click **OK** to close the dialog box.
- 6 Right-click the newly named resource and select **Bring this resource online**.
- 7 If you created more than one disk group for the application, repeat this procedure to add another Volume Manager Disk Group resource for another disk group.

Installing the application on the cluster nodes

The application program files must be installed on the same local drive of all the cluster nodes. The application data and log files or other files related to the application data are installed on the shared storage.

Refer to the Microsoft documentation for any specific requirements for the application in a failover cluster environment.

Pointers 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 install on both nodes at once.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Remember not to accept the default location for the application data and log files. Instead, browse to the dynamic volumes that were prepared previously.

Pointers for installing the application on the second node

- In Failover Cluster Management, 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**. The Drive Letter and Paths window appears. You can choose from the following:
 - To add a drive letter, click **Add**. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
 - To change a drive letter, click **Modify**. The **Assign a drive letter** drop-down list becomes available. Change the drive letter and click **OK**.

- To add a mount point, click **Add**, 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, and click **OK** to mount the volume.
- To change a mount point, you must remove it and then select the **Add** option to add it back. To remove it, select it in the Drive Letter and Paths window and click **Remove**.

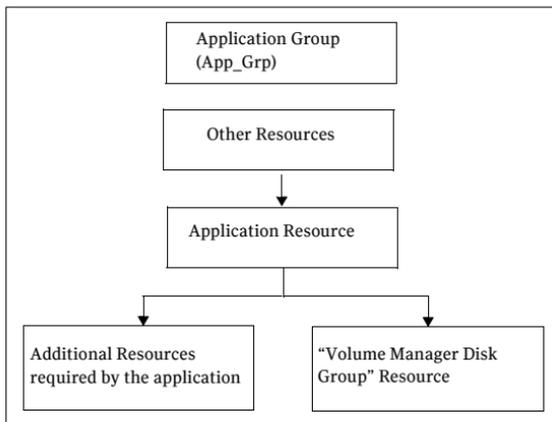
Completing the setup of the application group in the cluster

To make the application group functional in Microsoft clustering, the application resource needs to be added, as well as any other resources that are associated with the application. Also, dependencies need to be established for the resources. This section presents a summary of the process for completing the application group setup.

- Before creating 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.
- Refer to the application documentation for help on creating its resource and additional resources that may be required. You may need to create an IP address resource and a network name resource in addition to the Volume Manager Disk Group resource that you created earlier.
Ensure that you select the appropriate disk group resource as the storage resource on which the application resource is dependent.

The following dependency chart indicates the dependencies that are established.

Figure 20-5 Application group dependencies



- **Testing:** After the application group is set up, test it by moving the cluster resources to another node and then move them back.

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

This chapter contains the following topics:

- [“Tasks for deploying SFW and VVR with MSCS \(Windows Server 2003\)”](#) on page 759
- [“Part 1: Setting up the cluster on the primary site”](#) on page 763
- [“Part 2: Setting up the cluster on the secondary site”](#) on page 781
- [“Part 3: Adding the VVR components for replication”](#) on page 784
- [“Part 4: Maintaining normal operations and recovery procedures”](#) on page 800

Tasks for deploying SFW and VVR with MSCS (Windows Server 2003)

This chapter provides the steps for setting up a disaster recovery (DR) solution, using SFW with an MSCS cluster and VVR in a new installation on Windows Server 2003.

For information on setting up SFW with Microsoft failover clustering and VVR in a new installation on Windows Server 2008, see:

[Chapter 22, “Deploying SFW and VVR with Microsoft failover clustering”](#) on page 805

The example describes a generic database application in order to present general recommendations that apply to applications in a DR solution.

The process for setting up and working with the SFW-MSCS-VVR disaster recovery solution has four main parts:

- [Part 1: Setting up the cluster on the primary site](#)
- [Part 2: Setting up the cluster on the secondary site](#)
- [Part 3: Adding the VVR components for replication](#)
- [Part 4: Maintaining normal operations and recovery procedures](#)

The steps for setting up the MSCS cluster that were described in the High Availability section of this guide are the basic foundation on which this disaster recovery solution is built (see [Chapter 17, “Deploying SFW with MSCS” on page 609](#)). The main differences in the process of setting up the cluster for a disaster recovery rather than for HA alone are that you need to make sure that the VVR option is selected during the SFW installation and to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes. In setting up the secondary site, the cluster process is similar.

Once the two clusters are set up, one at the primary site and the other at the secondary site, VVR is used to enable replication from the primary site to the secondary site.

[Table 21-1](#) outlines the process for this configuration in more detail.

The high-level objectives and the tasks to complete each objective for the configuration are as follows:

Table 21-1 Task list for deploying SFW with MSCS and VVR

Objective	Tasks
“Reviewing the prerequisites and the configuration” on page 763	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites. ■ Review configuration requirements.
Part 1: Setting up the cluster on the primary site	
“Installing and configuring the hardware” on page 767	<ul style="list-style-type: none"> ■ Set up and configure the hardware according to the manufacturers’ instructions.
“Installing Windows and configuring network settings” on page 767	<ul style="list-style-type: none"> ■ Install the operating system on both nodes. ■ Make necessary networking settings on both nodes.
“Establishing the cluster under MSCS (primary site)” on page 769	<ul style="list-style-type: none"> ■ Refer to Microsoft documentation for instructions on establishing a cluster under MSCS.
“Installing SFW (primary site)” on page 769	<ul style="list-style-type: none"> ■ In the Options screen of the installer, select the VVR option.

Table 21-1 Task list for deploying SFW with MSCS and VVR (Continued)

Objective	Tasks
“Installing Veritas Volume Replicator Security Services (VxSAS)” on page 770	<ul style="list-style-type: none">■ Complete the steps to configure VxSAS.
“Creating SFW disk groups and volumes” on page 772	<ul style="list-style-type: none">■ In SFW on the primary cluster node, 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.■ The disk group for the quorum can be created later, if desired.
“Setting up a group for the application in MSCS” on page 773	<ul style="list-style-type: none">■ Create a group within MSCS for the application.■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application (Primary site)” on page 775	<ul style="list-style-type: none">■ Install the application program files on the local drive of the first node.■ Install files relating to the data and logs on the shared storage.■ Use Move Group to move the cluster resources to the second node.■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node.■ Install the application on the second node.
“Completing the setup of the application group in MSCS” on page 776	<ul style="list-style-type: none">■ Refer to the application documentation for help on creating its resource.■ Establish the appropriate dependencies.■ Test the application group by using the Move Group command to move the cluster resources to the other node.
“Changing the quorum resource to the dynamic mirrored quorum resource” on page 780	<ul style="list-style-type: none">■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.■ Make the disk group a Volume Manager Disk Group type resource in the default Cluster Group.■ Change the quorum resource to the dynamic mirrored quorum resource.

Table 21-1 Task list for deploying SFW with MSCS and VVR (Continued)

Objective	Tasks
“Testing of the cluster on the primary site” on page 781	<ul style="list-style-type: none"> ■ Use the Move Group command to move the cluster resources to the second node. Move them back to the first node. ■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.
<p>Part 2: Setting up the cluster on the secondary site</p>	
“Repeating cluster configuration steps for the secondary site” on page 783	<p>The tasks are:</p> <ul style="list-style-type: none"> ■ Installing and configuring hardware ■ Installing Windows and configuring network settings ■ Establishing the cluster under MSCS ■ Installing SFW ■ Installing Veritas Volume Replicator Security Services (VxSAS) ■ Creating SFW disk groups and volumes ■ Setting up a group for the application in MSCS ■ Installing the application on cluster nodes ■ Completing the setup of the MSCS application group ■ Changing the quorum resource to the dynamic quorum resource ■ Final testing of the cluster
<p>Part 3: Adding the VVR components for replication</p>	
“Configuring the Replicator Log volumes for VVR” on page 785	<ul style="list-style-type: none"> ■ Use SFW to create Replicator Log volumes for the primary and secondary sites.
“Setting up the Replicated Data Sets (RDS) for VVR” on page 787	<ul style="list-style-type: none"> ■ Create Replicated Data Sets with VVR’s Replicated Data Set wizard and start replication for the primary and secondary sites.
“Creating an RVG resource and setting the dependencies” on page 797	<ul style="list-style-type: none"> ■ In MSCS Cluster Administrator, create an RVG resource for replication. ■ Set the application resource dependency on the RVG resource. ■ Remove the direct dependency of the application resource on the Volume Manager Disk Group resource, VMDg.

Table 21-1 Task list for deploying SFW with MSCS and VVR (Continued)

Objective	Tasks
Part 4: Maintaining normal operations and recovery procedures	
“Normal operations: Monitoring the status of the replication” on page 800	<ul style="list-style-type: none">■ Monitor replication.■ Perform planned migration.
“Disaster recovery procedures” on page 801	<ul style="list-style-type: none">■ Complete the recovery procedures after the primary site goes down.

Part 1: Setting up the cluster on the primary site

This section provides more information on the steps for creating the cluster on the primary site.

Reviewing the prerequisites and the configuration

This topic describes the hardware and software requirements and gives an overview of the configuration.

Note: Before configuring the cluster, refer to the Microsoft documentation for MSCS requirements and the application documentation for application-specific 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 SFW with 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

Your system must run one of the following operating systems to install the SFW server software.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

The following table shows supported Windows Server 2003 operating systems.

Table 21-5 Supported Windows Server 2003 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2003	x86	Web Edition	SP2 required
Windows Server 2003	x86, x64, IA64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition, Datacenter Edition	SP2 required
Windows Server 2003 R2	x86, x64	Small Business Server	SP2 required
Windows Storage Server 2003 R2	x86, x64	Standard Edition, Enterprise Edition	SP2 required
Windows Unified Data Storage Server 2003	x86, x64	Standard Edition, Enterprise Edition	SP2 required

System requirements

- One CD-ROM drive accessible to each system on which you are installing MSCS.
- 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 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.

- Each system requires 1 GB of RAM for SFW.
- 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 Domain 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 (NICs), three for each server (two each for the private network and one for the public network) are required. You also need a static IP address for the cluster itself. Thus, you need seven IP addresses for each cluster.
- In addition, you need two more IP addresses for replication, one for the primary site and one for the secondary site.

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 21-6 Disk space requirements

Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

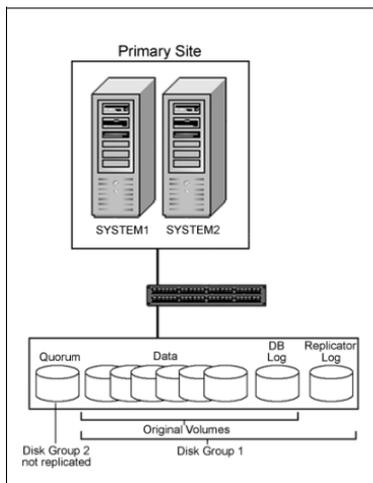
Reviewing the configuration

This configuration overview highlights the active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive

configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the configuration that includes both sites, see the illustration in the section [“About a disaster recovery solution”](#) on page 474.

Figure 21-1 DR configuration primary site



In addition, other items may be needed for this configuration:

- An MSCS cluster must be running to install SFW. Thus, you need to set up the hardware and install the operating system and MSCS on both systems in each cluster and establish the MSCS cluster before installing SFW.

Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Thus, Symantec recommends that you use a “rolling install” procedure to install SFW first on the inactive cluster node, then move the active 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*.

- SFW adds the advantage of the dynamic mirrored quorum. The main advantage of using SFW instead of MSCS alone is that MSCS by itself does not support mirroring the quorum resource; thus, with MSCS alone, the quorum disk itself is a point of failure for the cluster. SFW provides a dynamic mirrored quorum resource for MSCS. After SFW is installed on the cluster nodes, the next task is to 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, but you may want to wait until after you have your application installed, running, and tested with the cluster to convert the original basic quorum disk to the dynamic quorum volume. By waiting until the end of the process to convert from the basic physical disk quorum to the dynamic mirrored volume, you can make sure that the application is working first with the cluster and then add the dynamic quorum volume. The quorum disk group on each site does not get replicated because each cluster has its own quorum.

Installing and configuring the hardware

Refer to the hardware documentation and Microsoft documentation for specific details of your hardware setup.

As a best practice, Microsoft recommends that you wait until after the cluster is established on the first node before connecting the second node to the storage array to avoid data corruption on the disks.

Installing Windows and configuring network settings

This topic summarizes the steps for installing the operating system and configuring the network settings. For specific details, refer to the Microsoft documentation.

To install Windows and configure network settings

- 1 Install the operating system and MSCS on both servers.
MSCS is automatically installed with Windows Server 2003.
- 2 Establish the network settings for the NICs and the domain on both servers.
You need to establish static IP addresses for all six NICs—two private NICs and one public NIC for each system.

Establishing the cluster under MSCS (primary site)

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.

Note: To prepare for a rolling installation of SFW on Node A, Symantec recommends that you leave the cluster resources on Node B at this point.

Installing SFW (primary site)

The procedure for adding SFW support to the cluster on the primary site involves the same installation steps that were described earlier in the chapter on setting up a cluster with SFW and MSCS with one important difference: that you select the VVR option from the product installer Options screen.

Refer to “[Installing SFW](#)” on page 618 for detailed steps.

Installing Veritas Volume Replicator Security Services (VxSAS)

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server environment, 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 Launch the VVR Security Service Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxssascfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name (domain)\account)	Enter the administrative account name.
Password	Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains

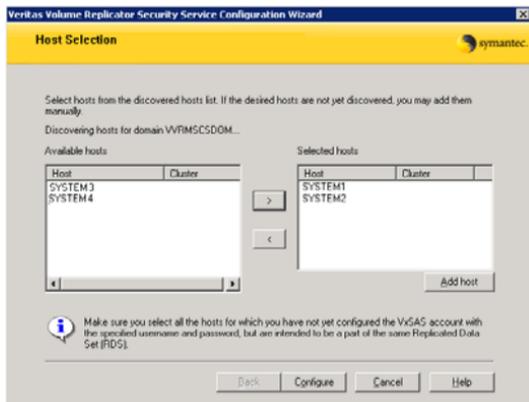
The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain

If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts

The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow keys. 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, click **Add host**. 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.

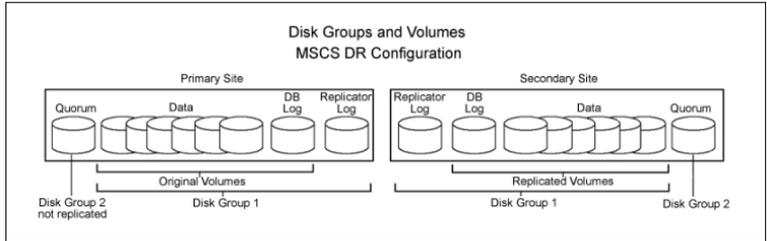
- 6 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- 7 Click **Finish** to exit the wizard.

Creating SFW disk groups and volumes

The following figure shows a typical setup of volumes for an MSCS VVR configuration with a database application. The example has one disk group for the application on each site.

If there are more application disk groups in your configuration, note that each disk group requires an additional Replicator Log volume. In the procedures described in this chapter, the Replicator Log volume will be created later; but you will need to allow sufficient disk space for the number of Replicator Log volumes required by your configuration. The quorum volume is not replicated to the second site and is in a separate disk group. It has to be created on each site and functions only on that site. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using three disks for the mirrored quorum for additional redundancy.

Figure 21-2 MSCS clustered database with disks for data, logs, and the quorum resource



Do not use the following types of volumes for the data and Replicator Log volumes; VVR does not support these types of volumes:

- SFW (software) RAID 5 volumes
- Volumes with commas in the names

For detailed steps in creating disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 626 in Chapter 8.

Setting up a group for the application in MSCS

Use Cluster Administrator in MSCS to set up a group for the application that will contain the SFW disk group or groups that were created for the application. The SFW disk groups are added to the MSCS application group as Volume Manager Disk Group (VMDG) resources.

After the application is installed on both nodes and its accompanying files are placed on the shared storage, complete the setup of the application group by adding the application itself as a resource and any other resources that are required. Dependencies need to be set between the resources in the group. Information on this task is included in [“Completing the setup of the application group in MSCS”](#) on page 776.

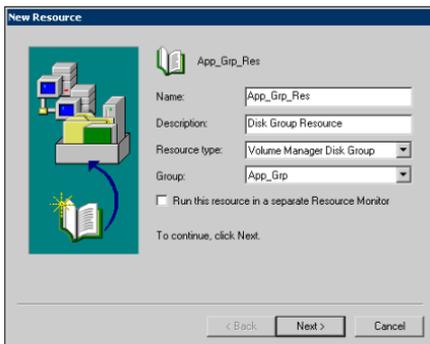
To set up the application group

- 1 Start Cluster Administrator in MSCS by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
Make sure you are connected to the appropriate cluster.

- 2 Create a new group by selecting the **Groups** node from the tree-view that is displayed on the left pane. Right-click to display the **Groups** context menu. Select **New > Group** from the menu.
The New Group window appears.
- 3 Specify a name for the group in the **Name** field (for example, App_Grp). Click **Next** to continue.
- 4 The Preferred Owners page appears. Make sure that all the preferred owners are added to the **Preferred Owners** list.
- 5 Click **Finish** to create the group.
You can now add resources to the group.

To add SFW disk groups as resources to the application group

- 1 Right-click on the MSCS group that you have created for the application and select **New > Resource**. The New Resource window appears.



- Specify a name for the disk group resource in the **Name** field.
- If required, you can add a description about the resource in the **Description** field. "Disk Group Resource" is an appropriate description.
- Specify the resource type by selecting **Volume Manager Disk Group** from the **Resource type** field drop-down list.

Note that the resource name has not been changed to Storage Foundation Disk Group.

- 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**.

The Possible Owners screen appears.

- 2 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.

The Dependencies screen appears.

- 3 On the Dependencies screen, click **Next**. You do not need to set any dependencies for a disk group resource.
- 4 In the next screen, make sure the appropriate SFW cluster disk group is selected from the drop-down list for the resource, and click **Finish** to complete the operation.

If there is more than one disk group for the application, you need to repeat the process. You can also add more resources at this time, as required by the application, or wait until after the application is installed. You will not be able to add the application resource until after the application is installed on both nodes.

Installing the application (Primary site)

This section gives more information on installing the application software on the two nodes of the primary site. Refer to the application documentation for any instructions on installing the application in a cluster.

Requirements

- The application program files need to be installed on the same local drive on all nodes. For example, if you install the application program files at **C:\Program Files\<application>** on one node, then these files must be installed at **C:\Program Files\<application>** on all the other nodes. Make sure that the same drive letter is available on all nodes and that there is adequate space for the installation.
- The data files and any associated files, such as log files, should be installed on the volumes under the clustered disk group or groups on the shared storage.
- To install the application on the second node, use the **Move Group** command to move the cluster resources to the second node.

Refer to [“Installing the application on cluster nodes”](#) on page 635 for more information on this task.

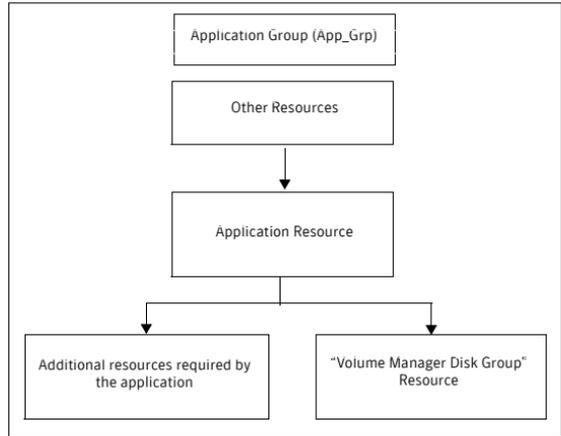
Completing the setup of the application group in MSCS

At this point in the process, additional steps make the application group functional in MSCS. The application resource is added, as well as any other resources that are associated with the application. Also, dependencies are established for the resources.

This section presents a high-level summary of the process for completing the application group setup.

- Before creating the application resource, make sure that all the other resources that you created—that is, the disk group resource and any additional application resources—are online.
- Refer to the application documentation for help on creating its resource and additional resources that may be required. You may need to create an IP Address resource and a Network Name resource in addition to the Volume Manager Disk Group resource you created earlier.
When creating the application resource, on the Dependencies screen select the **Volume Manager Disk Group** resource from “Available Resources” and add it to “Resource Dependencies.”
- The following dependency chart indicates the Dependencies that are established.

Application group dependencies



- **Testing:** After the application group is set up, test it by using the **Move Group** command to move the cluster resources to another node and then move them back.

Changing the quorum resource to a dynamic quorum resource

One of the key advantages of using SFW with MSCS to create a mirrored quorum resource to add fault tolerance to the quorum, thus protecting the cluster from failure, if the disk with the quorum fails. In the following procedure, transfer the cluster's quorum resource from a physical disk resource to a mirrored dynamic quorum resource. The tasks involved are:

- [“Creating a dynamic cluster disk group for the quorum with a mirrored volume”](#) on page 778
- [“Making the quorum cluster disk group an MSCS resource”](#) on page 778
- [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 780

Creating a dynamic cluster disk group for the quorum with a mirrored volume

If you have not completed this step earlier, use SFW to create a dynamic disk group for the quorum disks. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using three small disks for the mirrored quorum for additional redundancy.

If possible, use small disks, because the disk group will be used only for the quorum volume, which Microsoft recommends to be 500 MB. To create a three-way mirrored volume in the New Volume wizard, select the **Concatenated** layout, enable the **Mirrored** checkbox, and specify three mirrors. For full details on creating cluster disk groups and volumes, see [“Creating SFW disk groups and volumes”](#) on page 626.

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.

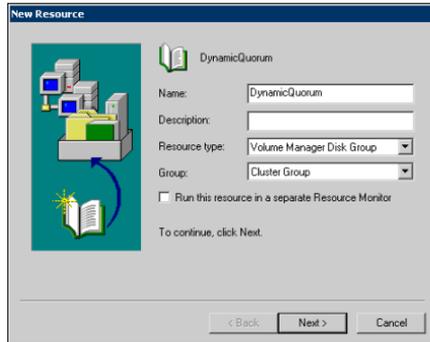
Making the quorum cluster disk group an MSCS resource

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 on that disk group and select **New > Resource**. The New Resource window appears.



- Specify a name for the disk group resource in the **Name** field, such as “DynamicQuorum.”
 - If necessary, you can 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.
Note that 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**.
- The Possible Owners screen appears.
- 3 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.
The Dependencies screen appears.
 - 4 On the Dependencies screen, click **Next**. You do not need to set any dependencies for a disk group resource.

- 5 In the next screen, 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 the prepared dynamic disk quorum resource.

To change the quorum resource to the dynamic mirrored quorum resource

- 1 Start Cluster Administrator in MSCS by selecting **Start > Settings > Control Panel > Administrative Tools > Cluster Administrator**.
- 2 From Cluster Administrator, right-click the cluster name in the tree view and 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**.

Testing of the cluster on the primary site

After the application is installed and the dynamic mirrored quorum is installed, test the cluster to make sure that it functions properly before adding the VVR components.

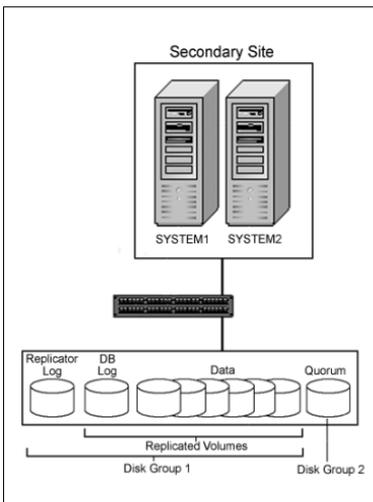
Refer to the section on testing an MSCS cluster, “[Verifying the cluster configuration](#)” on page 640.

Part 2: Setting up the cluster on the secondary site

On the secondary site, repeat the tasks performed on the primary site to create a cluster that duplicates the primary site’s disk groups and volumes.

The secondary disk groups and volumes should have the same names as those on the primary site. The data volumes should be the same sizes as the corresponding data volumes on the primary site. The log volume on the secondary site can be a different size, but Symantec recommends that the sizes be the same. Install the application on the secondary cluster nodes the same way as on the primary cluster.

Figure 21-3 DR configuration secondary site



Repeating cluster configuration steps for the secondary site

Refer to the guidelines provided earlier in this chapter to complete the same tasks on the secondary site prior to the application installation.

- [“Installing and configuring the hardware”](#) on page 767
- [“Installing Windows and configuring network settings”](#) on page 767
- [“Establishing the cluster under MSCS \(primary site\)”](#) on page 769
- [“Installing SFW \(primary site\)”](#) on page 769
- [“Installing Veritas Volume Replicator Security Services \(VxSAS\)”](#) on page 770
- [“Creating SFW disk groups and volumes”](#) on page 772
- [“Setting up a group for the application in MSCS”](#) on page 773

After completing these tasks, continue with the next section for instructions on installing the application on the secondary site.

Installing the application (Secondary site)

Installing the application on the secondary site is similar to installing it on the primary site.

See the section [“Installing the application \(Primary site\)”](#) on page 775.

Note: Before installing the application on the secondary site, offline all the resources in the MSCS application group on the primary site, except the disk group resource.

Completing the setup of the MSCS application group

Refer to [“Completing the setup of the application group in MSCS”](#) on page 776.

Changing the quorum resource to the dynamic quorum resource

Refer to [“Changing the quorum resource to the dynamic mirrored quorum resource”](#) on page 780.

Final testing of the cluster

Refer to [“Testing of the cluster on the primary site”](#) on page 781.

Before configuring VVR components

After both clusters are running, one on the primary site and one on the secondary site, you can add VVR components to the configuration.

Note: Before configuring the VVR components, on the secondary site offline all the resources in the application group, except the disk group resource.

Part 3: Adding the VVR components for replication

This section provides information on configuring the VVR components for replication. Topics include:

- [“VVR components overview”](#) on page 784
- [“Configuring the Replicator Log volumes for VVR”](#) on page 785
- [“Setting up the Replicated Data Sets \(RDS\) for VVR”](#) on page 787
- [“Creating an RVG resource and setting the dependencies”](#) on page 797

VVR components overview

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	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.
-----------------------	---

Configuring the Replicator Log volumes for VVR

Create the volume for the Replicator Log at each site. The task of creating the logs can also be done during the RDS creation process, but some storage administrators may prefer to do it manually (as is being done here) as a preparatory step to setting up the RDS.

Note: Before configuring the Replicator Log volumes, make sure that all the resources in the MSCS application group are offline, except the disk group resource. This task must be done on the primary site as well as the secondary site.

Note: To improve write performance, Symantec recommends that you create the Replicator Log volume on a different disk from the disks used for your application data volumes.

To configure the Replicator Log volumes for VVR

- 1 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.
- 2 Create a volume for the disk group that contains the storage group data:
 - On the System configuration tree, click the disk group where the log volume will be created (*Hostname>Disk Groups>Diskgroupname*).
 - Right-click on a disk group that has the volumes to be replicated, and click **New Volume**.
- 3 On the Welcome page of the New Volume wizard, click **Next**.
- 4 Select the disks for the volume:
 - Select the group name.

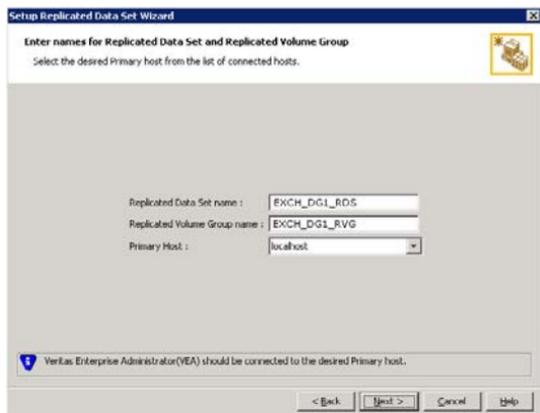
- Select **Manually select disks**.
 - Click the disk name.
 - Click **Add**.
 - After selecting all the necessary disks, click **Next**.
- 5 Specify the parameters of the volume:
- Enter the volume name.
 - Enter the size.
- The size of the Replicator Log volume varies for different environments. To determine the appropriate size for your environment, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.
- Select the volume layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 6 On the Add Drive Letter and Paths dialog box:
- Click **Do not assign a drive letter**.
 - Click **Next**.
- 7 When prompted to format the volume:
- Deselect **Format this volume**.
 - Click **Next**.
- 8 Click **Finish** to create the new volume.
- 9 If necessary, repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for any additional RVGs on the primary site.
- 10 Repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for additional disk groups on the secondary site.

Setting up the Replicated Data Sets (RDS) for VVR

Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

To create the Replicated Data Set

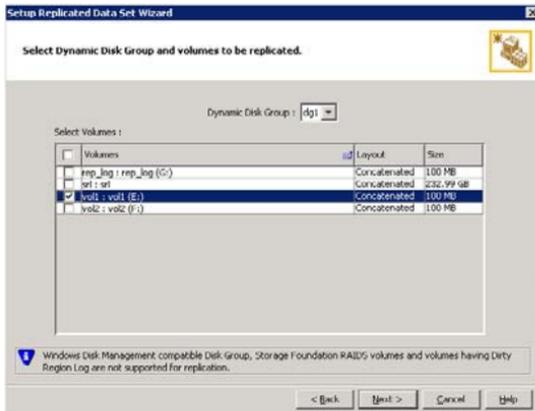
- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the 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 (APP_REPL_LOG). 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:

Setup Replicated Data Set Wizard

Edit replication settings
Edit replication settings or click next.

Primary side IP: 10.217.53.214

Secondary side IP: 10.217.53.215

Replication Mode: Synchronous Override

Replicator Log Protection: AutoDCM

Primary RLINK Name: Pri_RLINK

Secondary RLINK Name: Sec_RLINK

Advanced

DHCP addresses are not supported by VVR.

< Back Next > Cancel Help

- 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 MISSING.</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.

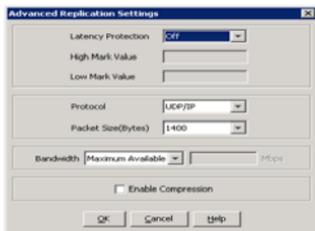
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

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 an RVG resource and setting the dependencies

This section describes additional tasks that must be done to complete the configuration of the MSCS application service group at both the primary and secondary sites. The tasks are:

- [Creating a replicated volume group \(RVG\) resource](#)
- [Setting the application resource dependency on the RVG resource](#)

Creating a replicated volume group (RVG) resource

To create an RVG resource

- 1 On the primary site, access Cluster Administrator by selecting **Start > Control Panel > Administrative Tools > Cluster Administrator**.
- 2 Right-click on the application cluster group that you created, and select **New > Resource**.
The New Resource screen appears.
 - Specify a name for the RVG resource in the **Name** field. For example, **RVGResource** is an appropriate name.
 - If desired, 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.
 - Configure a separate resource monitor process for the RVG resource by selecting the **Run this resource in a separate Resource Monitor** checkbox.
 - Click **Next**.
The Possible Owners screen appears.
- 3 By default, all the nodes in the cluster are listed as possible owners. Click **Next**.
The Dependencies screen appears.
 - 4 On the Dependencies screen, select the IP resource created earlier for VVR and the Volume Manager Disk Group resource from the “Available Resources” list in the left pane and add it to “Resource Dependencies” list in the right pane of the screen. Click **Next**.
 - 5 In the Replicated Volume Group Parameters screen, select the disk group for the RVG resource. Click **Finish**.
 - 6 Bring the RVG resource online.
 - 7 Repeat these same steps on the secondary site.

Setting the application resource dependency on the RVG resource

The application resource has a direct dependency on the Volume Manager Disk Group resource. With the addition of the RVG resource to the application group, the application’s dependency will change. The application will have a direct dependency on the RVG resource, which in turn depends on the Volume Manager Disk Group resource.

Note: The Volume Manager Disk Group resource represents the cluster disk groups created and managed by SFW.

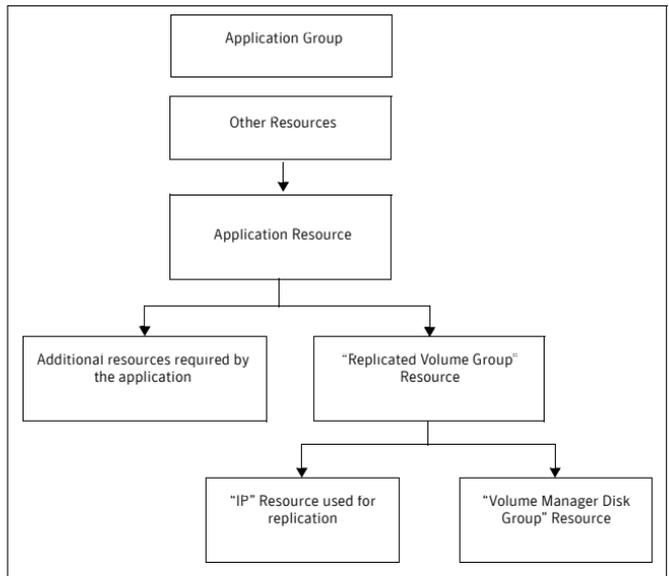
To set the application resource dependency on the RVG resource

- 1 Make sure the application resource is off-line before attempting to modify the dependencies.
- 2 Right-click on the application resource and select the **Properties > Dependencies** tab. This will display the Dependencies screen.
- 3 Click **Modify**.

- 4 Select the **Replicated Volume Group** resource from the “Available Resources” list and move it to the “Resource Dependencies” list.
- 5 Remove the **Volume Manager Disk Group** resource from the “Resource Dependencies” list. Click **OK**.
There is no longer a direct dependency between the application resource and the Volume Manager Disk Group resource.
- 6 The cluster configuration is now complete. Online the entire MSCS application group on the primary cluster.

The dependency chart that follows indicates the dependencies that have been established.

Figure 21-4 Dependencies of VVR-related resources



The chart shows only the VVR-related resources. Normally, there would be other resources involved in any clustered application. The main point of the

chart is to show that the RVG resource is now dependent on the Volume Manager Disk Group resource and the VVR virtual IP resource. The dependencies relationship has changed. The application resource is no longer directly dependent on the Volume Manager Disk Group resource.

Part 4: Maintaining normal operations and recovery procedures

This section provides tasks during normal operations of this solution and also describes the recovery process.

Normal operations: Monitoring the status of the replication

Under normal operating conditions, you can monitor the status of the replication using:

- The VEA GUI
- The command line interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Transfer the primary role to the host at the secondary site by 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.
- Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- Bring the application resource online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

Note: Any changes that you make to the data on the new primary will get replicated to the original primary, which is now the secondary.

Disaster recovery procedures

This section provides information on bringing up an application server on the secondary host in the event of a disaster. It also explains how to migrate the primary role back to the original primary host once it is returned to normal functioning after a disaster.

Bringing up the application on the secondary host

To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, 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 primary. Make sure that these drive letters are the same as those of the original primary.
- 5 Bring the application resource online.
Now you can start using the application on the new primary.

Restoring the primary host

After a disaster, if the original primary becomes available again, you may want to revert the role of the primary back to this host.

To restore the primary host

- 1 Depending on whether you performed **Takeover** with or without the **fast-failback** option, do one of the following:
 - For Takeover with the Fast-failback option:
The original primary, after it has recovered, will be in the **Acting as secondary** state. If the original primary is not in the **Acting as secondary** state, verify whether your network connection has been restored.
To synchronize this original primary and the new primary, use the **Resynchronize Secondaries** option from new primary's context menu.

- For Takeover without the Fast-failback option:
After performing a takeover without fast-failback, you must convert the original primary to a secondary by using the **Make Secondary** option.

Note: 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 the **Synchronize Automatically** option.
- 2 After the synchronization 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** from the menu that appears.
 - 3 Make sure that the volumes have retained the same drive letters as they had before the disaster.
 - 4 Bring the application resource online on the original primary.

Deploying SFW and VVR with Microsoft failover clustering

This chapter covers the following topics:

- [“Tasks for deploying SFW and VVR with Microsoft failover clustering \(Windows Server 2008\)”](#) on page 805
- [“Part 1: Setting up the cluster on the primary site”](#) on page 810
- [“Completing the primary site configuration”](#) on page 820
- [“Part 2: Setting up the cluster on the secondary site”](#) on page 821
- [“Part 3: Adding the VVR components for replication”](#) on page 823
- [“Part 4: Maintaining normal operations and recovery procedures”](#) on page 840

Tasks for deploying SFW and VVR with Microsoft failover clustering (Windows Server 2008)

You can set up a disaster recovery (DR) solution using SFW with a Microsoft failover cluster and VVR on Windows Server 2008.

For information on deploying SFW and VVR with MSCS on Windows Server 2003, see the following:

[Chapter 21, “Deploying SFW and VVR with MSCS” on page 759](#)

The example describes a generic database application in order to present general recommendations that apply to applications in a DR solution.

The process for setting up and working with the SFW-Microsoft failover cluster-VVR disaster recovery solution has four main parts:

- [Part 1: Setting up the cluster on the primary site](#)
- [Part 2: Setting up the cluster on the secondary site](#)
- [Part 3: Adding the VVR components for replication](#)
- [Part 4: Maintaining normal operations and recovery procedures](#)

The steps for setting up the failover cluster that were described in the High Availability section of this guide are the basic foundation on which this disaster recovery solution is built (see [Chapter 18, “Deploying SFW with Microsoft failover clustering” on page 643](#)).

The main differences in the process of setting up the cluster for a disaster recovery rather than for HA alone are that you need to make sure that the VVR option is selected during the SFW installation and to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes. In setting up the secondary site, the cluster process is similar.

Once the two clusters are set up, one at the primary site and the other at the secondary site, VVR is used to enable replication from the primary site to the secondary site.

The high-level objectives and the tasks to complete each objective for the configuration are as follows:

Table 22-1 Tasks for deploying SFW with Microsoft failover clustering and VVR

Objective	Tasks
“Reviewing the prerequisites and the configuration” on page 810	<ul style="list-style-type: none"> ■ Verify hardware and software prerequisites. ■ Review configuration requirements.
Part 1: Setting up the cluster on the primary site	
“Installing and configuring the hardware” on page 814	<ul style="list-style-type: none"> ■ Set up and configure the hardware according to the manufacturers’ instructions.
“Installing Windows and configuring network settings” on page 814	<ul style="list-style-type: none"> ■ Install the operating system on both nodes. ■ Make necessary networking settings on both nodes.
“Establishing the Microsoft failover cluster (primary site)” on page 814	<ul style="list-style-type: none"> ■ Refer to Microsoft documentation for instructions on establishing a Microsoft failover cluster.

Table 22-1 Tasks for deploying SFW with Microsoft failover clustering and VVR

(Continued)

Objective	Tasks
“Installing SFW (primary site)” on page 816	<ul style="list-style-type: none"> ■ In the Options screen of the installer, select the VVR option in addition to the Microsoft Cluster option.
“Installing Veritas Volume Replicator Security Services (VxSAS)” on page 816	<ul style="list-style-type: none"> ■ Complete the steps to configure VxSAS.
“Creating SFW disk groups and volumes” on page 819	<ul style="list-style-type: none"> ■ In SFW on the primary cluster node, 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. ■ The disk group for the quorum can be created later, if desired.
“Creating a group for the application in the failover cluster” on page 667	<ul style="list-style-type: none"> ■ Create a group for the application using the Microsoft Failover Cluster Management tool. ■ Include the cluster disk group or groups for the application as Volume Manager Disk Group type resources in the group.
“Installing the application on cluster nodes” on page 668	<ul style="list-style-type: none"> ■ Install the application program files on the local drive of the first node. ■ Install files relating to the data and logs on the shared storage. ■ Move the cluster resources to the second node. ■ Make sure that the volumes on the second node have the same drive letters or mount points that they had on the first node. ■ Install the application on the second node.
“Completing the setup of the application group in the failover cluster” on page 670	<ul style="list-style-type: none"> ■ Refer to the application documentation for help on creating its resource. ■ Establish the appropriate dependencies. ■ Test the application group by moving the cluster resources to the other node.

Table 22-1 Tasks for deploying SFW with Microsoft failover clustering and VVR

(Continued)

Objective	Tasks
<p>“Implementing a dynamic quorum resource” on page 671</p>	<ul style="list-style-type: none"> ■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier. ■ Make the disk group a Volume Manager Disk Group type resource in the default Cluster Group. ■ Change the quorum resource to the dynamic mirrored quorum resource.
<p>“Verifying the cluster configuration” on page 673</p>	<ul style="list-style-type: none"> ■ Move the cluster resources to the second node. Move them back to the first node. ■ Optionally, simulate a failure by turning off the power to the server that has control of the cluster resources.
<p>Part 2: Setting up the cluster on the secondary site</p>	
<p>“Repeating cluster configuration steps for the secondary site” on page 822</p>	<p>The tasks are:</p> <ul style="list-style-type: none"> ■ Installing and configuring hardware ■ Installing Windows and configuring network settings ■ Establishing the Microsoft failover cluster ■ Installing SFW with the VVR option and Microsoft Cluster option ■ Installing Veritas Volume Replicator Security Services (VxSAS) ■ Creating SFW disk groups and volumes ■ Setting up a group for the application in Failover Cluster Management ■ Installing the application on cluster nodes ■ Completing the setup of the application group in Failover Cluster Management ■ Changing the quorum resource to the dynamic quorum resource ■ Final testing of the cluster
<p>Part 3: Adding the VVR components for replication</p>	

Table 22-1 Tasks for deploying SFW with Microsoft failover clustering and VVR

(Continued)

Objective	Tasks
“Configuring the Replicator Log volumes for VVR” on page 824	<ul style="list-style-type: none"> ■ Use SFW to create Replicator Log volumes for the primary and secondary sites if this wasn't done earlier when configuring the SFW disk groups and volumes for the application.
“Setting up the Replicated Data Sets (RDS) for VVR” on page 826	<ul style="list-style-type: none"> ■ Create Replicated Data Sets with VVR's Replicated Data Set wizard and start replication for the primary and secondary sites.
“Creating resources for VVR” on page 836	<ul style="list-style-type: none"> ■ In Failover Cluster Management, create the network name and IP resource to be used for VVR replication.
“Creating an RVG resource and setting the dependencies” on page 837	<ul style="list-style-type: none"> ■ In Failover Cluster Management, create an RVG resource for replication. ■ Set the application resource dependency on the RVG resource. ■ Remove the direct dependency of the application resource on the Volume Manager Disk Group resource.
Part 4: Maintaining normal operations and recovery procedures	
“Normal operations: Monitoring the status of the replication” on page 840	<ul style="list-style-type: none"> ■ Monitor replication. ■ Perform planned migration.
“Disaster recovery procedures” on page 841	<ul style="list-style-type: none"> ■ Complete the recovery procedures after the primary site goes down.

Part 1: Setting up the cluster on the primary site

This section provides more information on the steps for creating the cluster on the primary site.

Reviewing the prerequisites and the configuration

This topic describes the hardware and software requirements and gives an overview of the configuration.

Note: Before configuring the cluster, refer to the Microsoft documentation for Microsoft failover cluster requirements and the application documentation for application-specific requirements.

Supported software for Microsoft failover clusters with SFW

The following software is supported for deploying SFW and Microsoft clustering on Windows Server 2008:

The following table shows supported Windows Server 2008 operating systems.

Note: SFW software for servers supports Hyper-V and parent partitions. SFW HA software for servers does not support Hyper-V and parent partitions.

Table 22-2 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows 2008 Server Core			
Windows 2008 Server Core			SP2
Windows 2008 R2 Server Core			
Windows Server 2008	x86,x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition, Small Business Server	SP1 or SP2 required
Windows Server 2008 for IA systems	IA64		SP1 or SP2

Table 22-2 Supported Windows Server 2008 operating systems

Windows Server	Platform	Edition	Service Pack
Windows Server 2008 R2	x64	Standard Edition, Enterprise Edition, Datacenter Edition, Web Edition	
Windows Server 2008 R2 for IA systems	IA64		
Windows Storage Server 2008			

System requirements

- One CD-ROM drive accessible to the system from 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 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.
- Each system requires a minimum of 1 GB of RAM.
- Each system in a 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 Domain 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 (NICs), three for each server (two each for the private network and one for the public network) are required. You also need a static IP address for the cluster itself. Thus, you need seven IP addresses for each cluster.
- In addition, you need two more IP addresses for replication, one for the primary site and one for the secondary site.

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 22-3 Disk space requirements

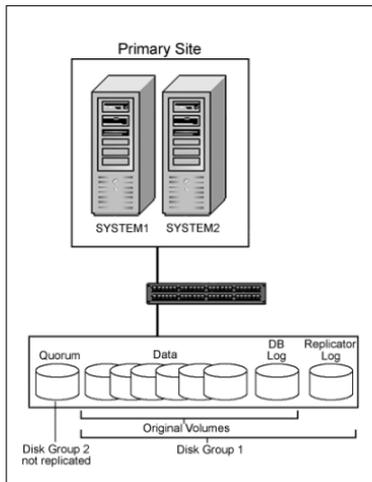
Installation options	Install directory/drive
SFW + all options + client components	1210 MB
SFW + all options	950 MB
Client components	354 MB

Reviewing the configuration

This configuration overview highlights the active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site, SYSTEM5 and SYSTEM6 on the secondary site), SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6. The figure that follows illustrates the cluster configuration on the primary site. For a view of the configuration that includes both sites, see the illustration in the section “[About a disaster recovery solution](#)” on page 474.

Figure 22-1 DR configuration primary site



This configuration does not include DMP. For information about DMP and clustering, see [“Adding DMP to a clustering configuration”](#) on page 67.

The following are some other key points about the configuration:

- A Microsoft failover cluster must be running before you install SFW. Installing SFW requires a reboot, but a reboot on the active cluster node causes it to fail over. Thus, Symantec recommends that you use a “rolling install” procedure to install SFW first on the inactive cluster node, then move the active cluster resources to the other node, and install on the now inactive node.
- SFW adds the advantage of the 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. For 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.

After SFW is installed on the cluster nodes, the next task is to 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 disk group and mirrored volume for the dynamic quorum resource.

The quorum disk group on each site does not get replicated because each cluster has its own quorum.

Installing and configuring the hardware

Refer to the hardware documentation and Microsoft documentation for specific details of your hardware setup.

As a best practice, Microsoft recommends that you wait until after the cluster is established on the first node before connecting the second node to the storage array in order to avoid corruption of data on the disks.

Installing Windows and configuring network settings

This topic summarizes the steps for installing the operating system and configuring the network settings. For specific details, refer to the Microsoft documentation.

To install Windows and configure network settings

- 1 Install the operating system and enable the Failover Clustering feature on both servers.
- 2 Establish the network settings for the NICs and the domain on both servers. You need to establish static IP addresses for all six NICs—two private NICs and one public NIC for each system.

Establishing the Microsoft failover cluster (primary site)

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**.

Part 1: Setting up the cluster on the primary site

- 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 (primary site)

The procedure for adding SFW support to the cluster on the primary site involves the same installation steps that were described earlier in the chapter on setting up a cluster with SFW and Microsoft failover clustering with one important difference: that you select the VVR option from the product installer Options screen.

See “[Installing SFW](#)” on page 653.

Installing Veritas Volume Replicator Security Services (VxSAS)

After installing SFW, you can configure the VxSAS service for Veritas Volume Replicator (VVR).

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server environment, 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 Launch the VVR Security Service Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxsascfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name Enter the administrative account name.
(domain)account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

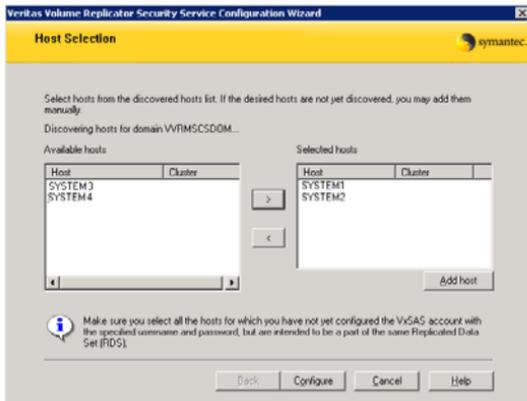
- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- 5 On the Host Selection panel, select the required hosts:



Selecting hosts The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. 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 displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with 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.

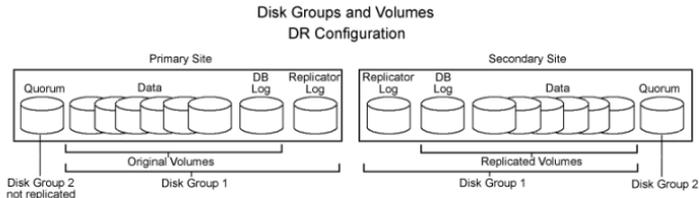
Creating SFW disk groups and volumes

The following figure shows a typical setup of volumes for an failover cluster with VVR configuration with a database application. The example has one disk group for the application on each site.

If there are more application disk groups in your configuration, note that each disk group requires an additional Replicator Log volume. In the procedures described in this chapter, the Replicator Log volume will be created later; but you will need to allow sufficient disk space for the number of Replicator Log volumes required by your configuration.

The quorum volume is not replicated to the second site and is in a separate disk group. It has to be created on each site and functions only on that site. The minimum number of disks for the mirrored quorum is two disks. Symantec recommends using three disks for the mirrored quorum for additional redundancy.

Figure 22-2 Microsoft clustered database with disks for data, logs, and the quorum resource



Do not use the following types of volumes for the data and Replicator Log volumes; VVR does not support these types of volumes:

- Storage Foundation for Windows (software) RAID 5 volumes
- Volumes with the Dirty Region Log (DRL)
- Volumes with a comma in their names

For the Replicator Log volume, in addition to the above types also make sure that the volume does not have a DCM.

For detailed steps in creating disk groups and volumes, see "[Creating SFW disk groups and volumes](#)" on page 659 in [Chapter 18, "Deploying SFW with Microsoft failover clustering"](#).

Completing the primary site configuration

The remainder of the tasks for the primary site configuration are identical to the tasks described in the chapter for configuring Microsoft failover clustering with SFW for high availability:

[Chapter 18, “Deploying SFW with Microsoft failover clustering”](#).

See the following topics in that chapter to complete configuring the primary site:

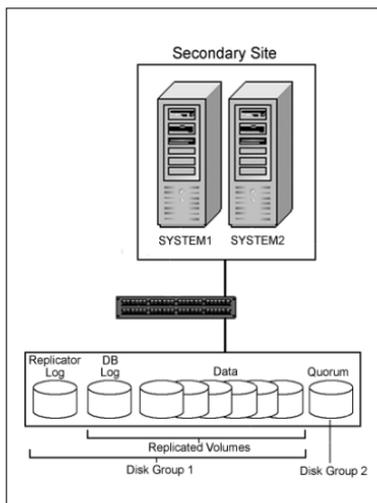
- [“Creating a group for the application in the failover cluster”](#) on page 667
- [“Installing the application on cluster nodes”](#) on page 668
- [“Completing the setup of the application group in the failover cluster”](#) on page 670
- [“Implementing a dynamic quorum resource”](#) on page 671
- [“Verifying the cluster configuration”](#) on page 673

Part 2: Setting up the cluster on the secondary site

On the secondary site, repeat the tasks performed on the primary site to create a cluster that duplicates the primary site's disk groups and volumes.

The secondary disk groups and volumes should have the same names as those on the primary site. The data volumes should be the same sizes as the corresponding data volumes on the primary site. The log volume on the secondary site can be a different size, but Symantec recommends that the sizes be the same. Install the application on the secondary cluster nodes the same as on the primary cluster.

Figure 22-3 DR configuration secondary site



Repeating cluster configuration steps for the secondary site

After configuring the cluster with SFW and the application on the primary site, follow the same procedures when configuring the secondary site. Follow the task list table.

See “[Tasks for deploying SFW and VVR with Microsoft failover clustering \(Windows Server 2008\)](#)” on page 805.

In addition, note the following special requirements for configuring the secondary site:

- During the creation of disk groups and volumes on 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
- Before installing the application on the secondary site, offline all the resources in the failover cluster application group on the primary site, except the Volume Manager Disk Group resource.

After both clusters are running, one on the primary site and one on the secondary site, you can add the VVR components to the configuration.

See “[Part 3: Adding the VVR components for replication](#)” on page 823.

Part 3: Adding the VVR components for replication

This section provides information on configuring the VVR components for replication. Topics include:

- “[VVR components overview](#)” on page 823
- “[Configuring the Replicator Log volumes for VVR](#)” on page 824
- “[Setting up the Replicated Data Sets \(RDS\) for VVR](#)” on page 826
- “[Creating resources for VVR](#)” on page 836
- “[Creating an RVG resource and setting the dependencies](#)” on page 837

VVR components overview

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>

Configuring the Replicator Log volumes for VVR

Note: Before configuring the Replicator Log volumes, make sure that all the resources in the failover cluster application group are offline, except the disk group resource. This task must be done on the primary site as well as the secondary site.

Create the volume for the Replicator Log at each site. The task of creating the logs can also be done during the RDS creation process, but some storage administrators may prefer to do it manually (as is being done here) as a preparatory step to setting up the RDS.

Note: To improve write performance, Symantec recommends that you create the Replicator Log volume on a different disk from the disks used for your application data volumes.

To configure the Replicator Log volumes for VVR

- 1 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.
- 2 Create a volume for the disk group that contains the storage group data:
 - On the System configuration tree, click the disk group where the log volume will be created (*Hostname>Disk Groups>Diskgroupname*).
 - Right-click on a disk group that has the volumes to be replicated, and click **New Volume**.
- 3 On the Welcome page of the New Volume wizard, click **Next**.
- 4 Select the disks for the volume:
 - Select the group name.
 - Select **Manually select disks**.
 - Click the disk name.
 - Click **Add**.
 - After selecting all the necessary disks, click **Next**.
- 5 Specify the parameters of the volume:
 - Enter the volume name.
 - Enter the size. The size of the Replicator Log volume varies for different environments. To determine the appropriate size for your

environment, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

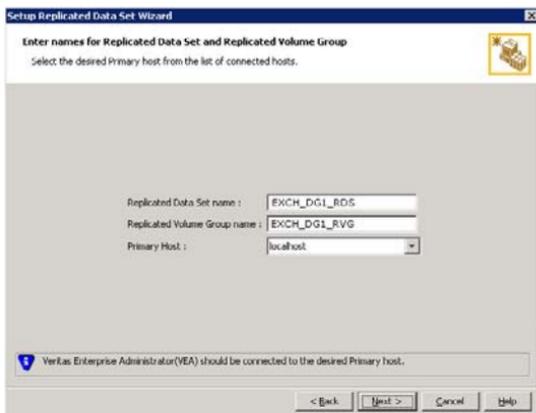
- Select the volume layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 6 On the Add Drive Letter and Paths dialog box:
 - Click **Do not assign a drive letter**.
 - Click **Next**.
 - 7 When prompted to format the volume:
 - Deselect **Format this volume**.
 - Click **Next**.
 - 8 Click **Finish** to create the new volume.
 - 9 If necessary, repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for any additional RVGs on the primary site.
 - 10 Repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for additional disk groups on the secondary site.

Setting up the Replicated Data Sets (RDS) for VVR

Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

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 (APP_REPL_LOG). 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:

Setup Replicated Data Set Wizard

Edit replication settings
Edit replication settings or click next.

Primary side IP: 10.217.53.214

Secondary side IP: 10.217.53.215

Replication Mode: Synchronous Override

Replicator Log Protection: AutoDCM

Primary RLINK Name: Pri_RLINK

Secondary RLINK Name: Sec_RLINK

Advanced

DHCP addresses are not supported by VVR.

< Back Next > Cancel Help

- 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. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. 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. 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 MISSING.</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

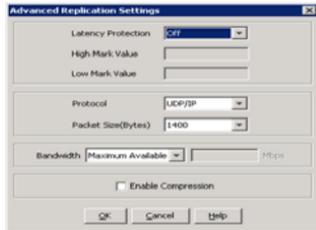
This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

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 resources for VVR

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.

Creating an RVG resource and setting the dependencies

This section describes additional tasks that must be done to complete the configuration of the Microsoft Failover Cluster application service group at both the primary and secondary sites. The tasks are:

- [Creating a replicated volume group \(RVG\) resource](#)
- [Setting the application resource dependency on the RVG resource](#)

Creating a replicated volume group (RVG) resource

To create a Replicated Volume Group (RVG) resource

- 1 In Failover Cluster Management, expand Services and Applications, right-click the application 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 application resource dependency on the RVG resource

When you specify resource dependencies, you control the order in which the cluster service brings resources online and takes them offline.

The application resource has a direct dependency on the Volume Manager Disk Group resource. With the addition of the RVG resource to the application group, the application's dependency will change. The application will have a direct dependency on the RVG resource, which in turn depends on the Volume Manager Disk Group resource.

Note: The Volume Manager Disk Group resource represents the cluster disk groups created and managed by SFW.

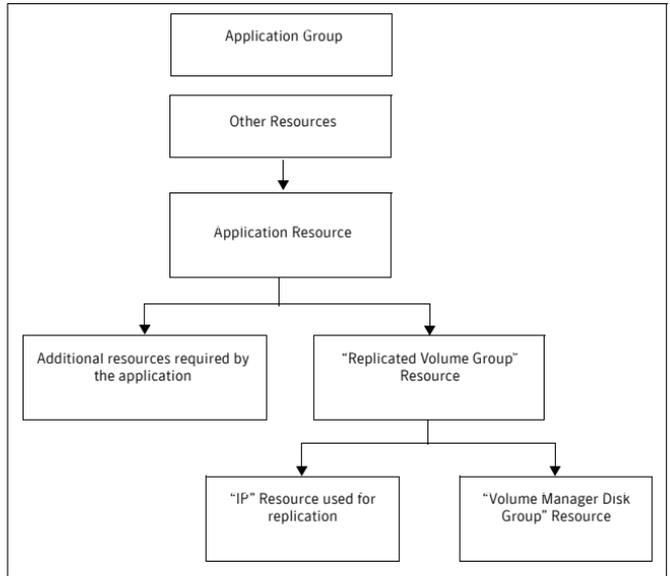
To set the application resource dependency on the RVG resource

- 1 Make sure the application resource is offline before attempting to modify the dependencies. Right-click the resource and click **Take this resource offline**.
- 2 Right-click the application resource and click **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.

The cluster configuration is now complete. Bring online the entire application group on the primary cluster.

The dependency chart that follows indicates the dependencies that have been established.

Figure 22-4 Dependencies of VVR-related resources



The chart shows only the VVR-related resources. Normally, there would be other resources involved in any clustered application. The main point of the chart is to show that the RVG resource is now dependent on the Volume Manager Disk Group resource and the VVR virtual IP resource. The dependencies relationship has changed. The application resource is no longer directly dependent on the Volume Manager Disk Group resource.

Part 4: Maintaining normal operations and recovery procedures

This section provides tasks during normal operations of this solution and also describes the recovery process.

Normal operations: Monitoring the status of the replication

Under normal operating conditions, you can monitor the status of the replication using:

- The VEA GUI
- The command line interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

To migrate the application to the secondary host

- 1 Take the Application resource offline on both the clusters. Stop the application so that volumes are not in use and secondary is up-to-date.
- 2 Transfer the primary role to the host at the secondary site by 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.
Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

Note: Any changes that you make to the data on the new primary will get replicated to the original primary, which is now the secondary.

Disaster recovery procedures

This section provides information on bringing up an application server on the secondary host in the event of a disaster. It also explains how to migrate the primary role back to the original Primary host once it is returned to normal functioning after a disaster.

Bringing up the application on the secondary host

To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, 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 Primary. Make sure that these drive letters are the same as those of the original Primary.
- 5 Bring the application resource online.

Now you can start using the application on the new Primary.

Restoring the primary host

After a disaster, if the original primary becomes available again, you may want to revert the role of the Primary back to this host.

To restore the primary host

- 1 Depending on whether you performed **Takeover** with or without the **fast-failback** option, do one of the following:

- For Takeover with the Fast-failback option:

The original primary, after it has recovered, will be in the `Acting as secondary` state. If the original Primary is not in the `Acting as secondary` state, verify whether your network connection has been restored.

To synchronize this original Primary and the new Primary, use the **Resynchronize Secondaries** option from new Primary's context menu.

- For Takeover without the Fast-failback option:

After performing a takeover without fast-failback, you must convert the original Primary to a Secondary by using the **Make Secondary** option.

Note: 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 the **Synchronize Automatically** option.

- 2 Take the application resource offline and stop the application.
- 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** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the application group online on the original Primary.

Server Consolidation

This section highlights server consolidation, the practice of consolidating server hardware, software, and data from multiple smaller servers to fewer, larger servers. This section also includes two sample configurations.

This section contains the following chapters:

- [Server consolidation overview](#)
- [Server consolidation configurations](#)

Server consolidation overview

This overview chapter describes server consolidation and focuses on how Veritas Storage Foundation for Windows (SFW) supports a server consolidation solution. The chapter's topics are:

- [“Server consolidation definition”](#) on page 847
- [“Need for implementing server consolidation”](#) on page 847
- [“Advantages of using SFW with server consolidation”](#) on page 848
- [“Overview of the server consolidation process”](#) on page 850

Server consolidation definition

Server consolidation is the consolidation of server hardware, applications, and data from multiple smaller, less powerful machines to fewer, more powerful servers. It involves sharing data in storage pools, usually in a storage area network (SAN).

Need for implementing server consolidation

Server consolidation provides the benefit of overall cost reduction by reducing the number of servers and their maintenance and administrative costs. Server consolidation also frees up space in the data center and improves security by reducing virus or software gateway risks, while improving service and availability. The larger, more powerful servers are better able to provide the computing power necessary to keep businesses competitive for the future.

Advantages of using SFW with server consolidation

Storage Foundation for Windows is ideally suited to support a server consolidation environment. Once servers are consolidated, SFW provides key features that assure fault tolerance and improve storage utilization. SFW's fault-tolerant features, such as software mirroring and RAID-5, Dynamic Multi-pathing (DMP), and clustering support assure high availability for consolidated storage, when business continuity is a requirement in a competitive business environment.

The SFW features that support server consolidation are:

- Ability to work in a heterogeneous storage environment
You are not tied to a solution offered by a single hardware vendor.
- Simple migration of data with disk group import and deport commands
If you have SFW disk groups already set up on multiple servers, you deport them on the source server, disconnect the attached storage, reattach the storage on the new larger server, and use the disk group import command to import the disk groups on the new server.
- Storage virtualization with software RAID volumes
Once the applications and data are consolidated on the new server, mirrored and RAID-5 volumes provide fault tolerance for critical data. Striped volumes add performance capabilities. Volumes that are both striped and mirrored offer both better performance and fault tolerance. Logical RAID volumes overcome the limitations of physical disks because these RAID volumes can span across disks and even disk arrays, thus assuring more efficient use of storage. Volumes can be configured online without restarting the server.
- Capacity management and online volume growth
Managing the space allocated for different functions is an important task that a system administrator must do on a consolidated server. SFW has a capacity monitoring function that alerts administrators when used space on a volume is near its capacity so that the volume can grow while it remains online. With this feature, you do not have to preallocate set amounts of storage for different purposes. More storage can be held in reserve in a pool for use only when it is needed. SFW volumes can be configured to increase capacity automatically when they pass a certain threshold.
- Online storage migration
If you need to take down a disk or even a whole disk array for maintenance, you can migrate the data online through the **Move Subdisk** command.
- Special features that support storage in a SAN

The importing and deporting of disk groups with host ID protection and private disk group protection can support storage in a SAN.

- **Dynamic Multi-pathing (DMP)**
The DMP software option increases performance of SAN-based disk arrays by spreading I/O between multiple paths to an array. Each path has a separate host adapter and cabling connecting the array and the server. If one path goes down, the DMP software automatically switches the storage associated with the failed path to an alternate path. Thus, the DMP software provides both fault tolerance for path failure and increases in performance through load balancing.
- **Clustering**
Storage Foundation for Windows supports clustering with MSCS and Storage Foundation HA for Windows includes Veritas Cluster Server. Clustering adds fault tolerance for servers. If one server in a clustered group of servers goes down, the storage of that server is taken over by another server in the cluster.
- **Additional fault tolerance features**
RAID-5 logging, dirty region logging, Hot Relocation, and FastResync (FR) increase the efficiency of the mirroring and RAID-5 functions in SFW.
- **Performance monitoring**
Online performance monitoring and tuning tools provide easy identification and minimization of I/O bottlenecks. These features allow you to increase throughput of the I/O in your system.

Overview of the server consolidation process

The server consolidation process involves more than just implementing the consolidation itself. It requires advance planning and approval of upper management. Here are some high-level steps:

- **Preliminary analysis:** Determine what servers need to be consolidated. Take into account the applications being used and the departments involved. Research the hardware and software needs and costs.
- **Design a plan for the consolidation and secure approval and budget from upper management.**
The primary justifications in the plan are cost savings and the need to remain competitive in today's business environment. The plan should also address IT management of the servers after the consolidation takes place.
- **Communicate with users about the proposed plan and identify the advantages of the plan before implementing the consolidation.**
Involve users in the planning process.
- **Do a proof of concept for the consolidation. Prototype the consolidation with a smaller number of servers that are not in production to see if your plan works.**
In the next section, two sample configurations are provided for demonstrating a proof of concept for consolidation.
- **Implement the consolidation on actual production servers.**
 - Purchase, install, and configure the new hardware and software for the migration.
 - Migrate the data.
 - Test to see that everything is working properly.
 - Put into effect new IT management processes for the consolidated servers.
- **In the months following the consolidation, implement a procedure to evaluate its effectiveness and the effectiveness of the IT management processes for the consolidated servers.**

Server consolidation configurations

The chapter's topics are:

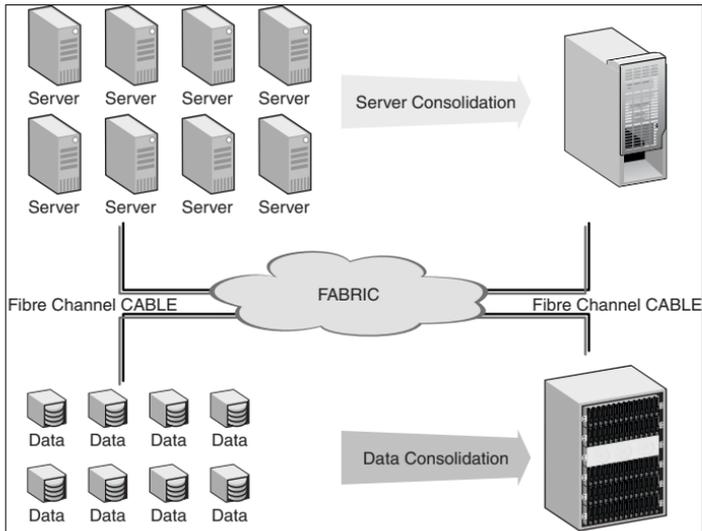
- [“Typical server consolidation configuration”](#) on page 852
- [“Server consolidation configuration 1—many to one”](#) on page 853
- [“Server consolidation configuration 2—many to two: Adding clustering and DMP”](#) on page 860
- [“SFW features that support server consolidation”](#) on page 868
- [“Server consolidation customer success story”](#) on page 869

Typical server consolidation configuration

This chapter provides two sample configurations that can be used as proof of concept for a consolidation.

The example shows a typical server consolidation situation. The consolidation could involve consolidating as many as 20 to 40 servers to one or two servers.

Figure 24-1 General server consolidation configuration



Proof of concept

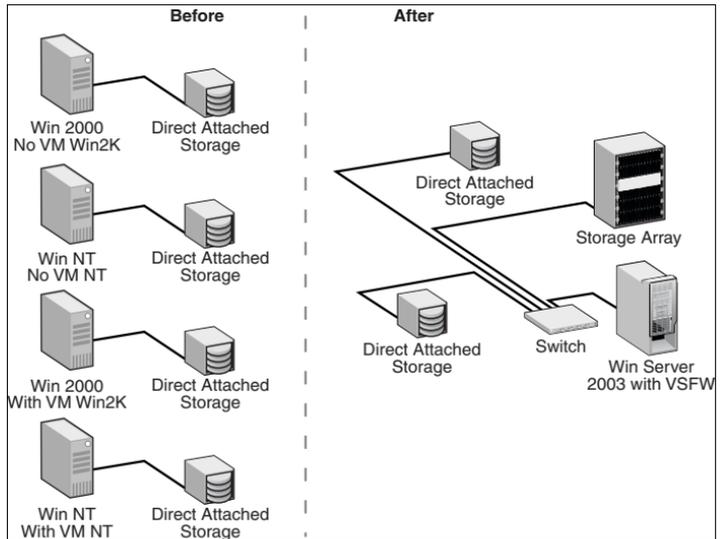
Testing the consolidation steps on a smaller number of servers provides an overview of the issues involved and how the process would work. In the configurations presented in this chapter, four servers are consolidated into one or two servers. The first configuration, which consolidates four smaller servers to one large server, provides fault tolerance through mirroring. In the second configuration, clustering and DMP are added to improve the fault tolerance, and an additional server is needed to support clustering. New, larger, more powerful

servers can be used in this proof of concept testing. Once the concept is tested, the main task is to migrate the data from the production servers to the new larger servers.

Server consolidation configuration 1—many to one

The following configuration illustrates consolidating many servers to one.

Figure 24-2 Proof of concept: Consolidating four small servers to one large server



About this configuration

In this configuration, four small servers are consolidated into a single larger server. The configuration also demonstrates that a server consolidation does not require that you eliminate all existing direct-attached storage units and replace them with large storage arrays. Setting up the storage on a SAN allows

you to use different combinations of storage devices and still derive the benefits from SFW's storage management features once the storage has been migrated from the small servers to a SAN.

Proof of concept

The four servers represent different Windows operating systems and Storage Foundation for Windows software combinations, which might be present in a production environment. The steps demonstrate that slightly different procedures are needed in preparing the storage for migration in each of these combinations.

In setting up your server consolidation configuration for proof of concept, select servers to migrate that have different combinations of typical hardware and software to determine the special requirements of such cases.

Phased approach: Flexible use of storage devices

In this example configuration, the steps are organized in phases:

- Preparing to consolidate
- Migrating the data to the large server
- Migrating data from the direct-attached storage to the storage array
- Adding the storage array
- Completing the consolidation process by migrating the storage from remaining servers

After the second phase in this example, all the direct-attached storage units have been detached from the small servers and are a storage pool on a SAN that is under the control of the new, large Windows Server 2003 system that is running SFW. You could stop at this point and still have many benefits from the storage that is now under SFW's management. If circumstances do not permit the purchase of a large storage array, you can simply use the existing direct-attached storage. Another alternative is to use both a storage array and some of the direct-attached storage. In this configuration and in Server Consolidation Configuration 2, using both a storage array and some of the direct-attached storage is shown.

The table below outlines the high-level objectives for implementing the configuration and the tasks for each objective:

Table 24-1 Tasks for server consolidation for many to one configuration

Objectives	Tasks
“Reviewing the configuration requirements” on page 856	<ul style="list-style-type: none">✓ Verify hardware and software requirements.
“Preparing to consolidate” on page 857	<ul style="list-style-type: none">✓ Make sure the data is backed up from the smaller servers before proceeding.✓ Set up the new large server and install the operating system and SFW. Connect it to the switch.✓ Prepare the data from each smaller server for consolidation by upgrading the server’s disks to dynamic disk groups, using either Disk Management or a version of Volume Manager for Windows.✓ Power down all the smaller servers and detach the storage.
“Migrating the data to the large server” on page 858	<ul style="list-style-type: none">✓ Reattach the direct-attached storage to the switch.✓ From the large server, import the disk groups from the direct-attached storage. The direct-attached storage is now attached to the SAN and is under the management of the large server that is running SFW. You could stop at this point if a large storage array is not available.
“Adding the storage array” on page 859	<ul style="list-style-type: none">✓ If you want to use a large storage array, set up the hardware array and connect it to the switch.✓ Migrate the data to the large storage array.

Table 24-1 Tasks for server consolidation for many to one configuration

Objectives	Tasks
“Completing the consolidation process” on page 859	✓ Migrate the data from the remaining servers.

Reviewing the configuration requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

Prerequisites

These procedures assume:

- Experience in setting up computer hardware, switches, and storage arrays
- Familiarity with the Windows operating systems and SFW/Volume Manager for Windows commands

Supported software

- Veritas Storage Foundation 5.1 for Windows (SFW)
or
Veritas Storage Foundation HA 5.1 for Windows (SFW HA)
- One of the following operating systems:
 - Windows Server 2003 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition (SP 1 required for all editions)
 - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
 - Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)
 - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
 - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

Hardware setup

- 4 smaller servers with direct-attached storage

- 1 larger, more powerful server
- Fibre switch and appropriate cabling for the SAN
- Hardware storage array

Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW or SFW HA.

Preparing to consolidate

In this phase, set up the large server and prepare the data for migration.

To prepare for consolidation

- 1 Identify the applications and data on the smaller servers that are a subset of the applications and data to be moved to the large server. You may want to have the users delete unnecessary files before the consolidation takes place.
- 2 Back up the data from the small servers.

Caution: back up the data from the small servers before proceeding.

- 3 Set up the large server and connect it to the switch.
- 4 Install the Windows Server 2003 operating system and Storage Foundation for Windows on the large server.
- 5 Prepare the data from each smaller server for migration by upgrading the server's disks to dynamic disk groups and powering down the server.

For Windows Server 2003 (no VM or SFW installed)

- Use Disk Management to upgrade basic disks to dynamic disks.
- Power down the server.

Migrating the data to the large server

Migrate the data to the large server. Perform the steps for each smaller server, one at a time.

To migrate the data to the large server

- 1 Disconnect the direct-attached storage from the small server.
- 2 Connect the direct-attached storage to the switch to make it accessible to the large server.

Note: All the direct-attached storage devices and the large server need to be in the same zone on the switch.

- 3 Using SFW on the large server, rescan the disks.
- 4 In SFW, import the disk groups from the direct-attached storage to make them a part of the storage that the large server manages.
Clear the host ID during the import process, if the source disk group was not created with SFW. A dialog box will come up for this purpose during the import command.
- 5 Assign drive letters to the imported disk groups.
On a Windows Server 2003 system, the default operating system setting requires the manual assignment of drive letters. Many administrators prefer to set drive letters manually rather than have the operating system do it.

Note: If you want the drive letters to be assigned automatically after a disk group is imported, use the `mountvol` command to change the default setting. Refer to the Microsoft documentation about the `mountvol` command for information on how to set up the automatic assignment of drive letters.

- 6 If desired, update the imported disk groups to the latest version of dynamic disk group type.
This is recommended to take advantage of the Windows Server 2003 features in SFW. Use the **Upgrade Dynamic Disk Group Version** command.
- 7 Test the data on the Windows Server 2003 system.

At this point, you can stop if you do not have a large storage array available. You can still take advantage of SFW's storage management features by having the direct-attached storage on the SAN. It is not necessary to have a large storage array to have these benefits.

Adding the storage array

If you have a large storage array available, the data may also be migrated to a hardware storage array on the SAN. You can eliminate all the direct-attached storage devices or keep them to increase your storage capacity. They can also be added into the configuration when needed.

To add the storage array

- 1 Set up and connect the hardware storage array to the switch.
On the switch, the hardware storage array must be in the same zone as the direct-attached storage devices and the large server.
- 2 Configure the array so that half of its disks are a mirror to the other half, using RAID-1. This provides fault tolerance to the storage.
- 3 Join the disk groups on the array storage and the direct-attached storage. This is done through the **Join Dynamic Disk Group** command.
- 4 Use the **Move Subdisk** command to move the volumes with data from the direct-attached storage to the array storage. You may want to keep some of the direct-attached storage on the SAN under the control of the large server. To access the **Move Subdisk** command:
 - Select the volume that contains the subdisk you want to move.
 - Click on the **Subdisks** tab in the right pane of the window.
 - Right-click the desired subdisk in the **Subdisks** tab and select **Move Subdisk** from the context menu.The **Move Subdisk** command also can be done by dragging and dropping the subdisks between disks in the Disk View. This method should be used with care to make sure that you do not move the subdisk to the wrong disk.
- 5 Test the data on the Windows Server 2003 system.
At this point, the migration of the storage from the four smaller servers is complete.

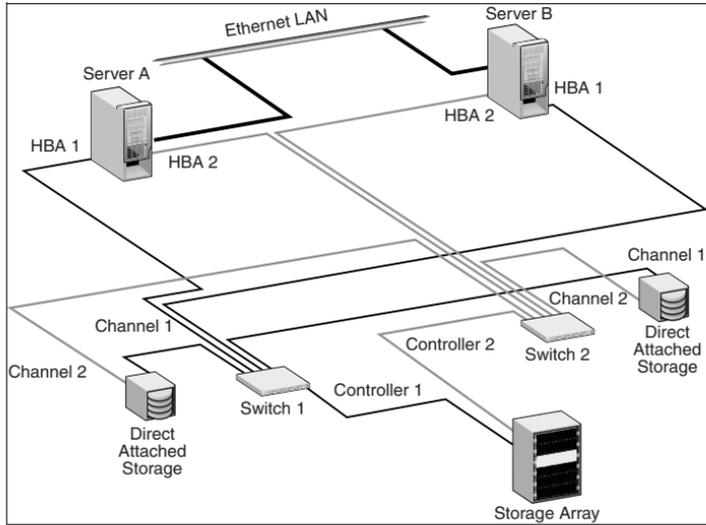
Completing the consolidation process

When you are satisfied that everything is working properly, migrate data from the remaining servers, using the methods shown in this configuration example.

Server consolidation configuration 2—many to two: Adding clustering and DMP

The following configuration consolidates many servers to two with MSCS clustering and DMP.

Figure 24-3 Adding fault tolerance with MSCS and DMP — requires two servers



About this configuration

This configuration is an upgrade to Server Consolidation Configuration 1, to add MSCS and DMP. Add a new server and host adapters, NICs, and a new switch.

The table below outlines the high-level objectives for implementing the configuration and the tasks for each objective:

Table 24-2 Tasks for server consolidation adding MSCS and DMP

Objectives	Tasks
“Reviewing the configuration requirements” on page 863	<ul style="list-style-type: none">✓ Verify hardware and software requirements.
“Adding the new hardware” on page 864	<ul style="list-style-type: none">✓ Add the new server, HBAs, network cards, and fibre switch.✓ Leave the second path for DMP unconnected on the existing server and the new server. It does not get connected until the end of the installation process.
“Establishing the MSCS cluster” on page 865	<ul style="list-style-type: none">✓ Refer to Microsoft instructions for establishing the cluster under MSCS.
“Adding SFW support to the cluster” on page 865	<ul style="list-style-type: none">✓ With Server B as the active cluster node, use Add or Remove Programs to add DMP and the MSCS support option to the first server.✓ With Server A as the active node, install SFW with the DMP and MSCS options to Server B.✓ Change the existing disk groups to cluster disk groups.✓ Prepare a disk group for the dynamic mirrored quorum.
“Setting up MSCS cluster groups for the applications” on page 866	<ul style="list-style-type: none">✓ If you have applications on the server that you want to cluster, create MSCS cluster groups for them.
“Installing applications on the second computer” on page 866	<ul style="list-style-type: none">✓ Install the applications’ program files on the local drive of Server B.

Table 24-2 Tasks for server consolidation adding MSCS and DMP

Objectives	Tasks
“Completing the setup of the application group in MSCS” on page 866	<ul style="list-style-type: none">■ Complete the cluster application group by adding resources and setting dependencies.
“Changing the quorum resource to the dynamic quorum resource” on page 867	<ul style="list-style-type: none">■ Create a dynamic disk group for the quorum with a mirrored volume if this task was not done earlier.■ Make that disk group a Volume Manager Disk Group type resource in the default Cluster Group.■ Change the quorum resource to the dynamic mirrored quorum resource.
“Verifying the cluster configuration” on page 867	<ul style="list-style-type: none">■ Test the cluster by moving the cluster resources to the other node.
“Enabling DMP” on page 867	<ul style="list-style-type: none">■ Using DMP, include the main storage array and, optionally, the direct-attached storage devices. Now attach the second path to the configuration and rescan.

More on DMP paths

In this configuration, there are two DMP paths, one going through Switch 1, which includes HBA 1 from Server A, HBA 1 from Server B, Channel 1 from the first direct-attached storage device, and Channel 1 from the second direct-attached storage device. The second path includes HBA 2 from Server A, HBA 2 from Server B, Channel 2 from the first direct-attached storage device, and Channel 2 from the second direct-attached storage device.

Caution: Do not have the second path to the storage connected to the SAN until DMP is installed and the storage array is included under DMP. If you allow two paths to the storage without DMP control, data can become corrupted.

The two switches keep the paths separate. You could use one large switch and zone it with two zones, one for each path.

In most DMP configurations, direct-attached storage is not included along with a storage array, but it is shown in this example to demonstrate that you can use direct-attached storage with DMP.

Note: When Storage Foundation for Windows is first installed, DMP control is not in effect. All arrays attached to the system come up as excluded. You must include the storage array and any direct-attached storage devices to enable DMP.

Reviewing the configuration requirements

Reviewing the prerequisites and the configuration allows you to gain an overall understanding of the configuration and its requirements.

Prerequisites

These procedures assume:

- Experience in setting up computer hardware, switches, and storage arrays
- Familiarity with the Windows operating systems and SFW commands

Supported software

- Veritas Storage Foundation 5.1 for Windows (SFW)
or
Veritas Storage Foundation HA 5.1 for Windows (SFW HA)
- One of the following operating systems:
 - Windows Server 2003 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition (SP 1 required for all editions)
 - Windows Server 2003 R2 (32-bit): Standard Edition, Enterprise Edition, or Datacenter Edition
 - Windows Server 2003 for 64-bit Itanium (IA64): Enterprise Edition or Datacenter Edition (SP 1 required for all editions)
 - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
 - Windows Server 2003 x64 Editions (for AMD 64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition

Hardware setup

Assume an original configuration of one large already consolidated server, one storage array, and two direct-attached storage devices that are all connected on a SAN.

Add:

- 1 large server of the same type as the first server
- Fibre switch
- 2 HBAs, one for each computer, required for DMP
- 6 network interface cards, 3 for each server (2 each for the private network and 1 for the public network)

Note: Refer to the Hardware Compatibility List on the Symantec Support web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW and SFW HA.

Recommendations

It is acceptable to use one NIC for the private network, but using two cards is strongly recommended to avoid making the private network a single point of failure in the configuration.

Refer to the Microsoft documentation for the specific requirements for the MSCS cluster. For example, you will need a static IP address for each network interface card and a static IP address for the cluster. The two clusters need to be members of the same domain.

Refer to [Chapter 17, “Deploying SFW with MSCS” on page 609](#) for more information on the process of setting up an MSCS cluster to work with SFW.

Adding the new hardware

Install the necessary hardware on both Server A and Server B.

To add the new hardware

- 1 Verify that your data from the large server is backed up before proceeding.
- 2 Install two host adapters in each server.

Caution: Do not connect the second path through HBA 2 on each server at this time.

- 3 Install the three network interface cards in each server. Do not make the connections between the two servers at this time.

- 4 Do any necessary configuration of the second switch without actually connecting it to the servers.

Establishing the MSCS cluster

Complete the steps necessary to install a cluster on Server A and Server B, using MSCS. Refer to the Microsoft documentation for the detailed instructions. The general steps are:

- Do the necessary network configuration steps on Server A.
- For example, establish the static IP addresses of the network cards and make sure a domain is set up that can be used by the two servers on the cluster.
- On Server A, access SFW and create a 500 MB partition on a disk that will be used as the quorum disk when the first node of the cluster is created. You may need to revert a dynamic disk to basic to implement this step.
- Create the first node of the cluster on Server A, using Cluster Administrator.
- Install the Windows Server 2003 operating system on Server B and do the networking configuration steps for Server B.
- Connect the networks between the two sites and verify their connectivity.
- Add the second node of the cluster to Server B.
- Test the cluster by moving the cluster resources from Server A to Server B. Server B becomes the active node. At this point, keep the control of the cluster with Server B.

Adding SFW support to the cluster

Use the following procedure to add SFW support to the cluster.

To add SFW support to the cluster

- 1 With the active node of the cluster on Server B, use **Add/Remove Programs** on Server A to add the MSCS and DMP options to SFW on that server and reboot. Then move the cluster resources back to server A. Server A is now the active node.

Note: If you reboot a server that has the active node of the cluster, it will fail over to the other node. You have more control of the situation by moving the resources to the other node before doing a reboot.

- 2 On Server B, install SFW with the DMP and MSCS options and reboot.

- 3 On Server A, which is now the active node of the cluster, use SFW to create a dynamic cluster disk group that will be used for the dynamic quorum. The disk group should contain three disks, and the disk size is recommended to be 500 MB. You need to create a three-way mirrored volume on the three disks with SFW. You can also use two disks, but three disks provide added redundancy.
- 4 Change the existing regular SFW dynamic disk groups on Server A to cluster disk groups.
A regular dynamic disk group is converted to a cluster disk group through the command line by using the command to import a disk group, `vxdg import`, with the `-s` option, the option that does the conversion. You will need to deport the disk groups first before you can import them. You can deport them through the GUI **Deport Dynamic Disk Group** command.

Setting up MSCS cluster groups for the applications

If you have applications on the server that you want to cluster, you need to set up an MSCS cluster group for each application. Set up the groups first before the application is installed because if the application is cluster-aware, it may need to reference the cluster group. For detailed steps on setting up MSCS cluster groups, see [“Setting up a group for the application in MSCS”](#) on page 633.

Note that you will not be able to finish setting up the resources for the group until the application is installed on the second node.

Installing applications on the second computer

If you have one or more applications on the existing computer and you want their data and associated files to be clustered, you need to install the applications on the local drive of the new computer. The applications may be cluster-aware and require specific procedures to install. Refer to the application documentation.

For tips on installing applications in an MSCS environment, see [“Installing the application on cluster nodes”](#) on page 635.

Completing the setup of the application group in MSCS

Once the application is installed, complete the configuration of the application group in MSCS. For details, see [“Completing the setup of the application group in MSCS”](#) on page 636.

Changing the quorum resource to the dynamic quorum resource

For details about changing the quorum resource to the dynamic quorum resource, see [“Implementing a dynamic quorum resource”](#) on page 637.

Verifying the cluster configuration

Verify that the cluster can fail over by moving the cluster group manually between the nodes to make sure it works properly. For details, see [“Verifying the cluster configuration”](#) on page 640.

Enabling DMP

These steps assume that SFW with the DMP option has been installed. See [“Adding SFW support to the cluster”](#) on page 865 in this example.

To enable DMP

- 1 With SFW on the first server, bring up DMP and include the disks on the storage array and optionally the two direct-attached storage devices. To include each storage array or direct-attached storage device under DMP control:
 - a Display the Array Settings screen for the device you are including by doing the following:
 - In the tree view under the **Disks** icon, select a disk from the storage array.
 - In the right pane, click the **Paths** tab for the disk. Only one path should display in the **Paths** tab, since the disk is not yet under DMP control.
 - Right-click the path and select **Array Settings** from the path context menu that comes up.
 - The Array Settings window comes up. The **Exclude** checkbox is checked.
 - b Uncheck the **Exclude** checkbox.
- 2 Using appropriate cables, connect the second path on Server A to Switch 2.
 - a Connect the path through Server A, HBA 2, Channel 2 of the direct-attached storage, and Controller 2 of the large storage array.
 - b Complete any necessary configuration of the switch.
- 3 Go to **Actions** and select **Rescan** to verify that two paths are shown under the **Paths** tab. This indicates that one set of disks has two paths and that DMP is installed correctly.

4 Complete [step 1](#) to [step 3](#) on Server B.

MSCS and DMP are now set up, and the upgraded configuration steps are complete.

SFW features that support server consolidation

With consolidated servers, Storage Foundation for Windows has multiple features that assure fault tolerance and improve storage utilization. Many of those features are highlighted in the section “[Advantages of using SFW with server consolidation](#)” on page 848.

The following section adds more information about some of the features. It describes how to create a script for Automatic Volume Growth based on capacity and gives a high-level view of SFW features for supporting storage in a SAN and for performance management. Topics in this section include:

- [Automatic volume growth](#)
- [Features that support storage in a SAN](#)
- [Performance monitoring](#)

Automatic volume growth

Storage Foundation for Windows comes with an Automatic Volume Growth feature that monitors the capacity of dynamic volumes and automatically increases the size of the volume when used space on it reaches a predetermined size.

With this procedure, you can conserve disk space on your servers because space is distributed automatically on an as-needed basis. You do not have to be available to allocate the additional disk space when it is required.

Features that support storage in a SAN

In a SAN environment, it is important to protect storage so that it cannot be accessed by more than one host at a time. SFW provides the feature of private dynamic disk group protection that protects a disk group with a SCSI reservation so that other hosts cannot access the data. For more information on this feature, see the *Veritas Storage Foundation Administrator's Guide*. Clustering is another way to protect the storage in a SAN. It also uses a SCSI reservation to keep the disk group from being accessed by other hosts in a SAN.

Performance monitoring

The statistics feature of SFW provides I/O statistics to allow performance tuning to improve overall disk and system performance. Through the Online Monitoring window, hot spots are identified. A hot spot is an area of high I/O activity that may cause bottlenecks in I/O throughput. If a disk has these hot spots, consider moving one or more of its subdisks to another disk that shows below-average I/O activity. For more information on this topic, see the *Veritas Storage Foundation Administrator's Guide*.

Server consolidation customer success story

At El Camino Hospital (ECH) in Mountain View, California—known as “The Hospital of Silicon Valley” and a world leader in the use of IT innovations and high-tech devices—there is never a good time for a problem to occur or to take systems offline for maintenance or upgrades. Even 99% uptime may not be good enough when patient care is at stake.

“We aim for 99.999%, because those clinical systems have to be there for physicians and nurses all the time,” said Joe Wagner, CTO of El Camino Hospital. “With Veritas, we deliver all departments, all applications, to all end users, 24 x 7. We’ve seen instant results and tangible cost savings, for a 24-month return on investment of about \$3.4 million.”

El Camino Hospital had been using Veritas backup products for years when it began researching software-based approaches that could cost-effectively deliver high availability. Like most hospitals, ECH was also looking for ways to consolidate servers to save money on hardware, software, and maintenance. “We came to the conclusion that Veritas’s high-availability solution could also deliver the server consolidation we were looking for,” said Wagner.

Veritas Consulting helped ECH design and implement a solution based on a many-to-one (many active to one standby) cluster strategy. “With Veritas, we have 15 live clusters or systems able to fail over to a single machine,” explains Wagner. “This means huge cost savings compared to hardware-based high-availability solutions.” The fact that Veritas is not platform-specific was also key to selecting this approach.

The return on the hospital’s investment has been “tremendous,” Wagner reported. “For more than a year now, we’ve been seeing cost savings in servers, storage, operating systems, network interface cards, host adapters, ports, power supplies, tape drives, and licenses, for a 24-month ROI of about \$3.4 million,” he explained. Significant savings have also been realized through cost avoidance and staff redeployments.

In addition to optimizing the ROI on the Veritas products the hospital had purchased, Veritas Consulting also worked closely with hospital staff to create a

comprehensive business continuity plan for an upcoming data center move, without disrupting the critical ongoing activity in the hospital. “Veritas has solved our immediate need for high availability, and now it’s enabling us to address the broader issues of business continuity planning,” said Wagner. “Veritas is our one-stop shop for continuous, cost-effective computing at El Camino Hospital.”

Appendix

- [Deploying Disaster Recovery: Manual implementation](#)

Deploying Disaster Recovery: Manual implementation

This chapter provides the steps for setting up a disaster recovery (DR) solution, using SFW HA with the Veritas Volume Replicator (VVR) and Global Cluster Option (GCO) in a new installation. The example describes a generic database application.

Note: This chapter covers the “manual” method of deploying disaster recovery. A newer method uses the Solutions Configuration Center and the Disaster Recovery (DR) wizard to clone storage configuration and service groups and set up replication. See [Chapter 14, “Deploying disaster recovery: New application installation” on page 481](#) .

For examples of the SFW HA disaster recovery solution with specific applications, see the other Solutions Guides included with this release: *Veritas Storage Foundation and High Availability Solutions Solutions Guide for Microsoft Exchange* and *Veritas Storage Foundation and High Availability Solutions Solutions Guide for Microsoft SQL*.

The process of setting up and working with the SFW-VVR disaster recovery solution has five parts:

- [“Part 1: Setting up the cluster on the primary site”](#) on page 879
- [“Part 2: Setting up the parallel environment on the secondary site”](#) on page 921
- [“Part 3: Adding the VVR components for replication”](#) on page 922
- [“Part 4: Adding GCO components for wide-area recovery”](#) on page 937

- [“Part 5: Maintaining: Normal Operations and recovery procedures”](#) on page 943

The following table outlines the DR process for this configuration in more detail.

Table A-1 Task list for deploying disaster recovery

Objective	Tasks
“Reviewing the configuration” on page 876	<ul style="list-style-type: none"> ■ Understanding the DR configuration
“Part 1: Setting up the cluster on the primary site.”	
“Installing SFW HA” on page 879	<ul style="list-style-type: none"> ■ Verify the driver signing option for the systems. ■ Install SFW HA <ul style="list-style-type: none"> ■ Select the option to install VVR. ■ Select the Global Cluster Option for VCS to enable wide-area failover.
“Configuring the VVR security service” on page 879	<ul style="list-style-type: none"> ■ Complete the steps to configure VxSAS
“Configuring the cluster” on page 882	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node. ■ Configure cluster components using the Veritas Cluster Server Configuration wizard. ■ Set up secure communication for the cluster.
“Configuring disk groups and volumes” on page 899	<ul style="list-style-type: none"> ■ Create disk groups. ■ Create volumes.

Table A-1 Task list for deploying disaster recovery (Continued)

Objective	Tasks
“Installing the application on cluster nodes” on page 906	<ul style="list-style-type: none"> ■ Install the application program files on the local drive of the first node. ■ Install files relating to the data and logs on the shared storage. ■ Deport the disk groups on the first node and import them on the second node. ■ Make sure that the volumes on the second node have the same drive letters or mount points as on the first node. ■ Install the application on the second node.
“Creating VCS service groups” on page 909	<ul style="list-style-type: none"> ■ Use an appropriate method to create the VCS service group for the application.
“Verifying the cluster configuration” on page 919	<ul style="list-style-type: none"> ■ Switch the service group to the second node. ■ Switch it back to the first node
“Part 2: Setting up the parallel environment on the secondary site.”	
	<ul style="list-style-type: none"> ■ Setting up a parallel environment on the secondary site involves: <ul style="list-style-type: none"> ■ Reviewing the requirements ■ Installing and configuring hardware ■ Configuring the network settings ■ Installing SFW HA ■ Configuring the VVR security service ■ Configuring disk groups and volumes ■ Configuring the cluster ■ Installing the application ■ Configuring the VCS service group for the application ■ Verifying the cluster configuration
“Part 3: Adding the VVR components for replication.”	

Table A-1 Task list for deploying disaster recovery (Continued)

Objective	Tasks
“Configuring the Replicator Log volumes for VVR” on page 923	<ul style="list-style-type: none"> ■ Use SFW to create Replicator Log volumes for the primary and secondary sites.
“Setting up the replicated data sets (RDS) for VVR” on page 925	<ul style="list-style-type: none"> ■ Create Replicated Data Sets with VVR’s Replicated Data Set wizard and start replication for the primary and secondary sites.
“Creating the VVR RVG Service group” on page 936	<ul style="list-style-type: none"> ■ Create a VVR RVG service group for the replicated volume group.
“Part 4: Adding GCO components for wide-area recovery.”	
“Linking clusters by adding a remote cluster” on page 938	<ul style="list-style-type: none"> ■ Create a global cluster by adding the first cluster to the second one through the command Add/Delete Remote Cluster.
“Converting a local Service group to a global group” on page 939	<ul style="list-style-type: none"> ■ Convert service groups that are common to all clusters to global service groups.
“Part 5: Maintaining: Normal Operations and recovery procedures.”	
“Normal operations: Monitoring the status of the replication” on page 943	<ul style="list-style-type: none"> ■ Monitor replication ■ Perform planned migration
“Disaster recovery procedures” on page 944	<ul style="list-style-type: none"> ■ Complete the recovery procedures after the primary site goes down.

Reviewing the configuration

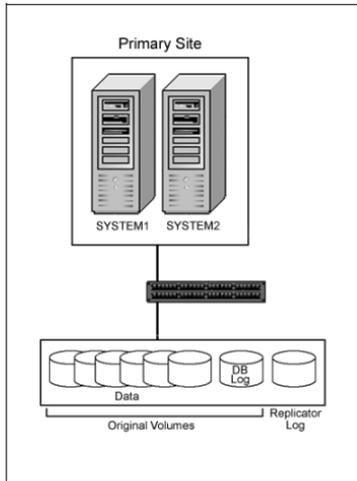
This configuration overview describes active/passive high availability within a cluster and disaster recovery between two sites. In an active/passive configuration, one or more application virtual servers can exist in a cluster, but each server must be managed by a service group configured with a distinct set of nodes in the cluster.

Active/passive clusters involve one-to-one failover capabilities. For instance, if you have two nodes on each site (SYSTEM1 and SYSTEM2 on the primary site,

SYSTEM5 and SYSTEM6 on the secondary site), then SYSTEM1 can fail over to SYSTEM2, and SYSTEM5 can fail over to SYSTEM6.

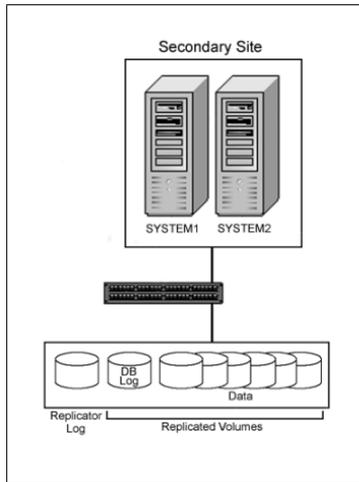
The figure that follows illustrates the cluster configuration on the primary site.

Figure A-1 DR configuration primary site



The following illustration shows the secondary site.

Figure A-2 DR configuration secondary site



Once the two clusters are set up, one at the primary site and the other at the secondary site, VVR is used to enable replication from the primary site to the secondary site.

The Global Cluster Option allows the two clusters to become global clusters and to be able to fail over to one another. Normally, two independent clusters cannot fail over to each other.

Part 1: Setting up the cluster on the primary site

This section details the steps for creating the cluster on the primary site.

The steps for setting up the cluster described in the High Availability section of this guide are the basic foundation for this disaster recovery solution. See [Chapter 8, “Deploying SFW HA for high availability: New installation” on page 77](#).

The main differences in the process of setting up the cluster for a disaster recovery, rather than for HA alone, are that you need to make sure that the VVR and the GCO options are selected during the SFW HA installation. You also need to configure the Veritas Volume Replicator Security Service (VxSAS) after the installation completes.

Installing SFW HA

To prepare for and install SFW HA on the primary site, use the following instructions from [Chapter 8, “Deploying SFW HA for high availability: New installation”](#):

- [“Reviewing the requirements” on page 80](#)
- [“Configuring the storage hardware and network” on page 88](#)
- [“Installing Veritas Storage Foundation and High Availability for Windows” on page 91](#)

Make sure that the VVR and the GCO options are selected during the SFW HA installation.

Configuring the VVR security service

If you are using VVR replication, you must configure the VxSAS service on all cluster nodes on both the primary and secondary sites. You can wait until you configure the secondary site and complete this procedure for both sites, or you can configure the primary site now and the secondary site later.

Complete the following procedure to configure the VxSAS service for VVR.

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. In a Windows Server environment, 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 Launch the VVR Security Service Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard**.
or
Type `vxsascfg.exe` at the command prompt.
- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows and then click **Next**:

Account name Enter the administrative account name.
(domain)\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

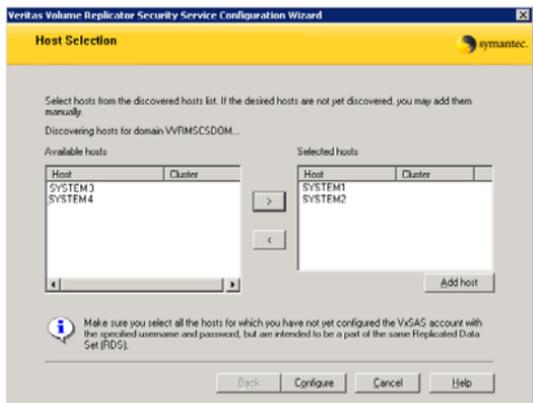
- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong and then click **Next**:

Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain If the domain name that you require is not displayed, click **Add domain**. This displays a dialog that allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

- On the Host Selection panel, select the required hosts:



Selecting hosts

The Available hosts pane lists the hosts that are present in the specified domain.

Move the appropriate host 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, click **Add host**. 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.

- After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- Click **Finish** to exit the wizard.

Configuring the cluster

After installing SFW HA using the installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters. The GCO option is needed for a disaster recovery solution. Thus, once you have completed the cluster configuration steps described in this section, complete the steps in the section that follows to configure the GCO option.

Complete the following tasks before creating a cluster:

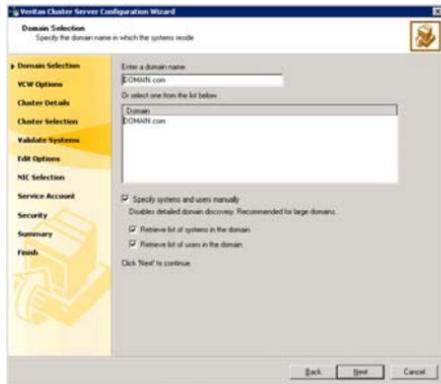
- Verify that each node uses a static IP address (DHCP is not supported) and that name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

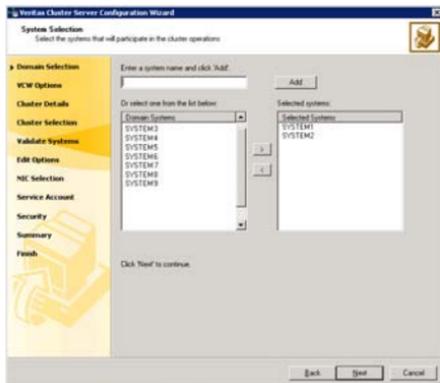
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 884.
 - To specify systems and user names manually (recommended for large domains):
 - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
 - Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 884. Otherwise proceed to the next step.
- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**.
Do not specify systems that are part of another cluster.
Proceed to [step 8](#) on page 884.

- 6 On the System Selection panel, specify the systems to form a cluster and then click **Next**. Do not select systems that are part of another cluster.



Enter the name of the system and click **Add** to add the system to the Selected Systems list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

- 7 The System Report panel displays the validation status, whether *Accepted* or *Rejected*, of all the systems you specified earlier. Review the status and then click **Next**.

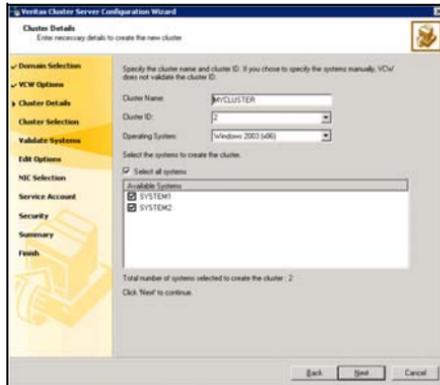
A system can be rejected for any of the following reasons:

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

Select a system to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



- Cluster Name** Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
- Cluster ID** Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.
- Caution:** If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.
- Operating System** From the drop-down list, select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

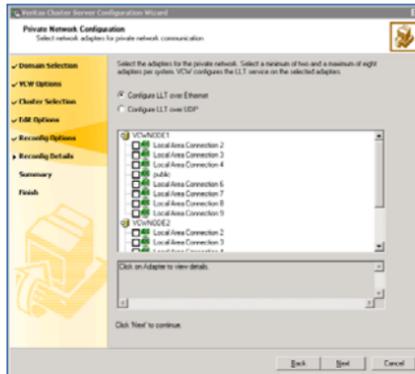
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 888.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer.

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:



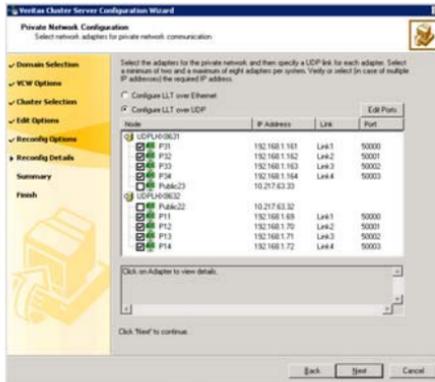
- Select the check boxes next to the two NICs to be assigned to the private network.

Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.

- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
- To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

The wizard configures the LLT service (over ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:



- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links.

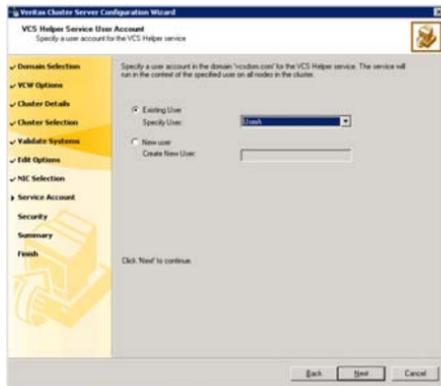
Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports will be used for the private network communication.

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

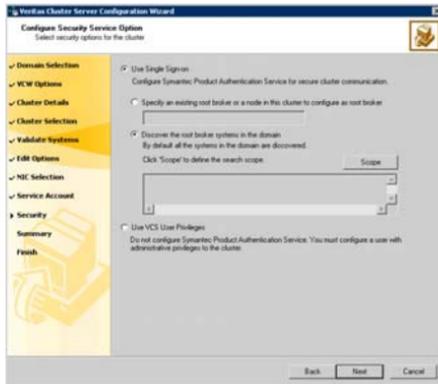
This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 883, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type ***2003*** in the adjacent field, click **Add** and then click **OK**.

Table A-2 contains some more examples of search criteria.

Table A-2 Search criteria examples

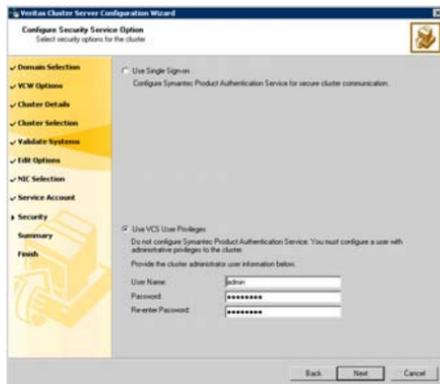
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

You can add multiple search criterion; the wizard will search for systems that match all the conditions specified.

- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

If the root broker is a cluster node, the wizard configures the other cluster nodes as authentication brokers. If the root broker is outside the cluster, the wizard configures all the cluster nodes as authentication brokers.

- To use VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**.

The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.

The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters.

To configure the ClusterService group later, click **Finish**.

At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

Note: After configuring the cluster you must not change the names of the nodes that are part of the cluster. If you wish to change a node name, run this wizard to remove the node from the cluster, rename the system, and then run this wizard again to add that system to the cluster.

- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** check box to configure notification of important events to designated recipients.
 - Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.

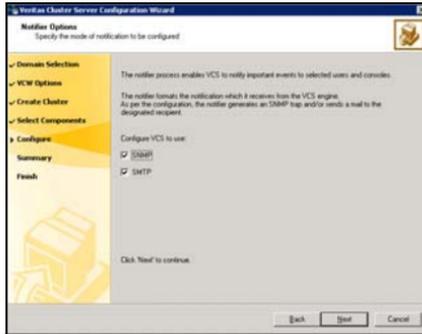
The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option

Configuring notification

This section describes steps to configure notification.

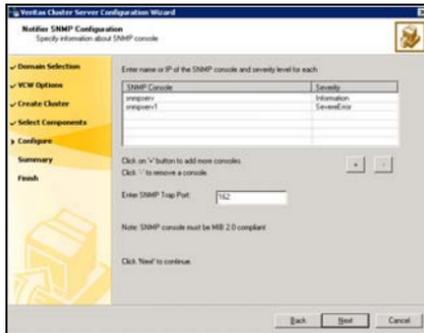
To configure notification

- 1 On the Notifier Options panel, specify the mode of notification to be configured and click **Next**.

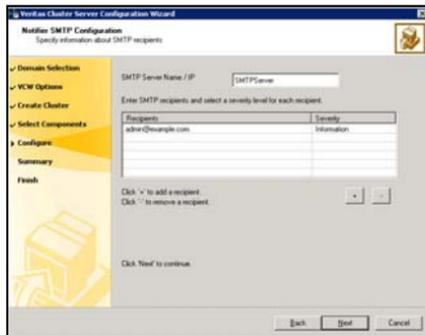


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

- 2 If you chose to configure SNMP, specify information about the SNMP console and click **Next**.

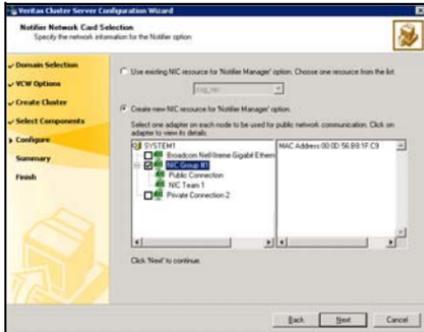


- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click + to add a field; click - to remove a field.
 - Enter an SNMP trap port. The default value is 162.
- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



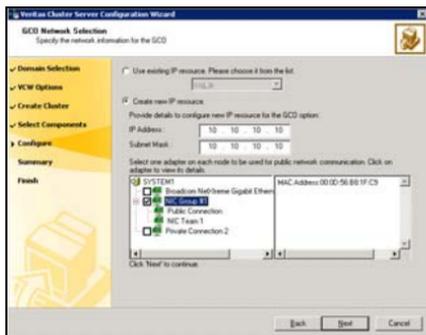
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.
 - 6 If you chose to configure global cluster components, proceed to [“Configuring the wide-area connector process for global clusters”](#) on page 897. Otherwise, click **Finish** to exit the wizard.

Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
 - 3 Click **Finish** to exit the wizard.

Configuring global cluster components

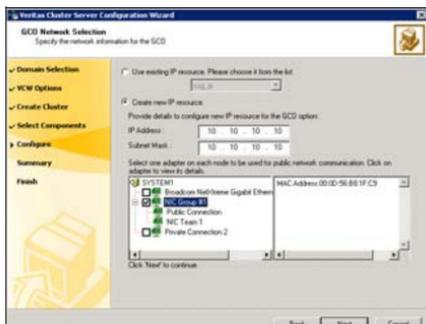
The next task is to identify an IP resource for the wide-area connector that is required for inter-cluster communication. If the cluster has a ClusterService

group configured, you can use the IP address configured in the service group or configure a new IP address.

This task does not set up a global cluster environment. That process is done later and is described in “[Part 4: Adding GCO components for wide-area recovery](#)” on page 937.

To configure an IP resource for GCO

- 1 In the GCO Network Selection panel, specify the network information:

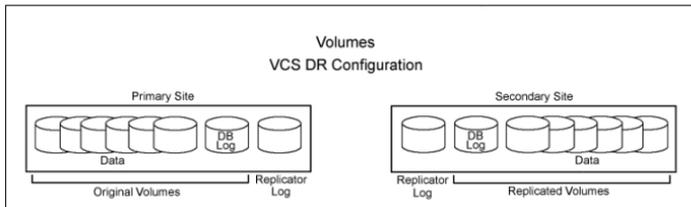


- If the cluster has a ClusterService group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask.
Make sure the specified IP address has a DNS entry.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
 - Click **Next**.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts.
 - 3 Click **Configure**.
 - 4 Click **Finish** to exit the wizard.

Configuring disk groups and volumes

The following figure shows a typical setup of volumes for a VCS disaster recovery configuration with a database application. The example has one disk group on each site.

Figure A-3 VCS clustered database volumes, DB log, and Replicator Log



Use Veritas Storage Foundation for Windows to create cluster disk groups and dynamic volumes for the application on the shared storage. A dynamic disk group is a collection of one or more disks that behave as a single storage repository. Within each disk group, you can have dynamic volumes with different RAID layouts.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - Control Panel - System Settings). See the Veritas Storage Foundation Administrator's Guide for more information.

Configuring disk groups and volumes involves the following tasks:

- "[Planning disk groups and volumes](#)" on page 900
- "[Configuring disk groups and volumes](#)" on page 899
- "[Creating dynamic volumes](#)" on page 903

Planning disk groups and volumes

Decide how you want to organize the disk groups and the number and type of volumes you want to create. Some considerations are:

- The number of disk groups that are needed
The number of disk groups depends on your application and the planned organization of the data. VCS requires that the application program files be installed on the local drive of the server. Data files and other related files, such as logs, are placed on the shared storage. Typically, a main organizational unit in your application, such as the storage group in Microsoft Exchange, would be contained in a single disk group.
- The type of volumes you want to create
 - Mirrored and RAID-5 volumes provide fault tolerance for critical data.
 - Striped volumes add performance capability.
 - Volumes that are both mirrored and striped offer both performance and fault tolerance.

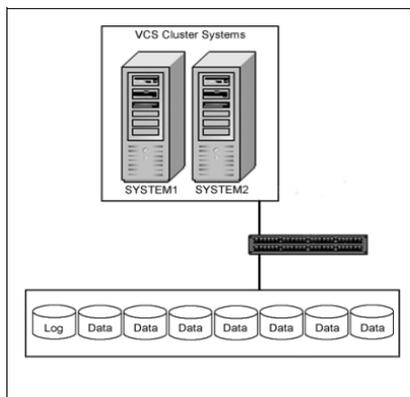
Note: If you plan to use replication software, such as VVR, do not use RAID-5 volumes. This does not apply to hardware RAID-5. VVR also does not support volumes with commas in the names.

Recommendations:

- Use mirrored volumes for logs.
- Use striped or mirrored striped volumes for data.
- The implications of backup and restore operations for the disk group setup.
- The sizes of databases and logs, which depend on the traffic load.

The following illustration shows a typical setup of disks for a clustered database application with shared storage. The log volume takes a single disk. Volumes for data and associated files take the remaining disks. Because you have dynamic volumes, the volumes can span multiple disks. You can have a mirrored striped volume that uses the disks other than the log disk.

Figure A-4 VCS clustered database with disks for data and the log



Creating dynamic cluster disk groups

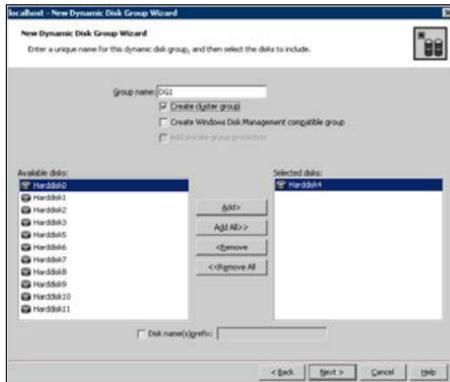
Follow the steps in this section to create one or more disk groups for your application.

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 disk group name (for example, DG1).
- 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.

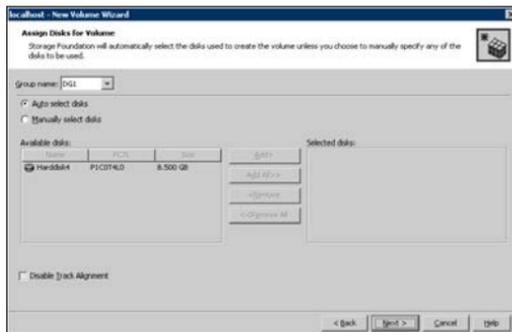
- 8 Click **Finish** to create the new disk group.

Creating dynamic volumes

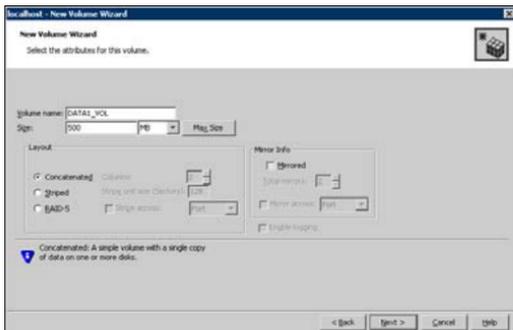
Once the disk groups are created, make the disks within them usable by creating the dynamic volumes that will store data.

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.
- 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.



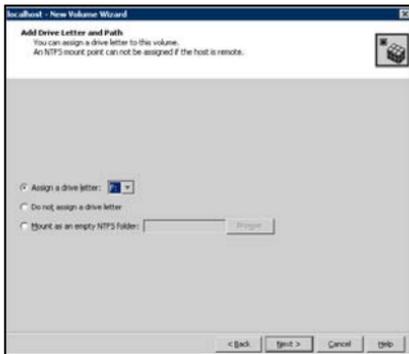
- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable 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

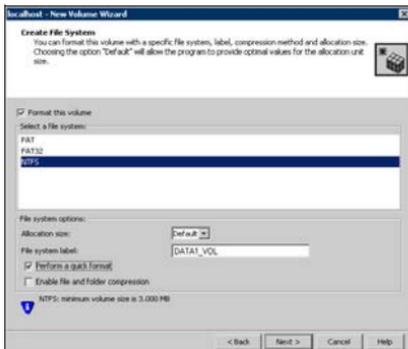
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- 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 additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

Installing the application on cluster nodes

VCS requires that the application program files be installed on the same local drive of all cluster nodes and that the application data and log files or other files related to the application data be installed on the shared storage.

Pointers 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.
- Make sure that the disk groups and volumes are imported and thus mounted on the server before you install the application.
- If you have just created the disk groups and volumes, they will be mounted and accessible. When a disk group is created, it is automatically imported on that node. You can verify that the disk group and volumes are accessible if you can see the disk group and volume icons in the VEA GUI for the server.
- All nodes of the clustered application need to share the same virtual name and IP address.
- Remember not to accept the default location for the application data and log files when installing the application. Instead, click to browse to the dynamic volumes that were prepared previously.

Pointers for installing the application on the second node

- To install the application on the second node, deport any disk groups from the first node and import them on the second node. Steps for deporting and importing disk groups are in the section [“Deporting and importing a disk group”](#) on page 907.
- You need to 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. To change a drive letter or mount point, see instructions in the section [“To add or change a drive letter or mount point”](#) on page 908.
- 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. You would then restart the service after the application is installed.

Deporting and importing a disk group

This topic describes the steps for deporting and importing a disk group in order to install the application on the second node.

To deport a disk group on the first node

- 1 If SFW is not already running, start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**).

- 2 Navigate to **dynamic disk groups** on the node on which the dynamic disk group is currently imported.
- 3 Right-click the dynamic disk group to be deported and click **Deport**.

To import the dynamic disk group on the second node

- 1 Start the Veritas Enterprise Administrator (**Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**).
- 2 Navigate to **dynamic disk groups** on the node to which you will import the dynamic disk group.
- 3 Right-click the dynamic disk group to be imported and click **Import**. There may be no drive letter associated with an existing dynamic volume when it is imported to a computer for the first time. Use SFW to add or change drive letters, as needed. Make sure that drive letters or mount points for the volumes on the second node are the same as were used on the first node.

To add or change a drive letter or mount point

- 1 In SFW, 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**. The Drive Letter and Paths window appears.
- 3 To add a drive letter, click the **Add** radio button. The **Assign a drive letter** drop-down list becomes available. Assign a drive letter and click **OK**.
- 4 To change a drive letter, click the **Modify** radio button. The **Assign a drive letter** drop-down list becomes available. Select the new drive letter and click **OK**.
- 5 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, and click **OK** to mount the volume.

A mount point is also referred to as a "drive path."

To change a mount point, you must remove it and recreate it ([step 5](#)). To remove it, select it in the Drive Letter and Paths window and click the **Remove** radio button.

Creating VCS service groups

This section on VCS service groups first describes the VCS service group and then presents an example of creating a service group with a generic database application.

About VCS service groups

In order for VCS to be able to monitor and fail over an application in a cluster, the application must be included in a VCS service group.

A service group is a collection of resources working together to provide application services to clients. It can also relate to a print or a file share that does not contain a specific application. A service group's resources fail over as a group to another cluster node when there is an application failure or server failure on the active node.

VCS has a collection of defined resource types. For each type, VCS has a corresponding agent that provides a type-specific logic to control resources. The bundled agents come with the program and apply to a group of resources that are commonly used with applications.

VCS provides multiple methods for creating a service group. For some applications and server roles, wizards are provided. In addition, you can use the VCS Java Console, as well as an Application Configuration wizard. A command line interface can be used to create a service group as well.

All these different methods accomplish the same purposes:

- Defining the cluster resources and their attributes.
- Setting their dependencies; for example, a NIC resource depends on an IP resource.
- Logically grouping the resources together.
- Providing capabilities for monitoring the service group and taking it online or offline.

For more information refer to the *Veritas Cluster Server Administrator's Guide*. The next section illustrates how to create a VCS service group using the Application Configuration wizard.

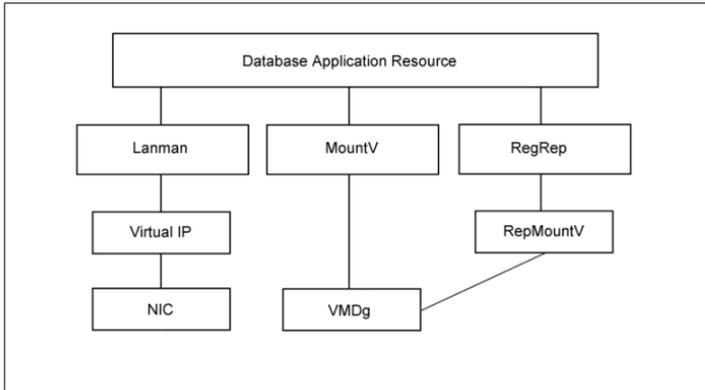
Service group example with a generic database application

The following steps show an example of creating a service group for a generic database application.

Application description

Assume that the application is a database application that runs as a Windows service. Its resources with their dependencies are shown in the chart that follows:

Figure A-5 Database application resources for the Service group example



- The Lanman resource makes the application available to clients. It depends on other resources that are associated with it.
- The virtual IP resource identifies the cluster and allows the cluster to communicate across the network. It depends on the NIC being configured for it to function.
- The MountV resource mounts the SFW disk group volumes and depends on the VMDg resource, which includes the SFW disk groups.
- The RegRep resource replicates the registry of the active cluster node and depends on the RepMountV resource and the VMDg resource.

The resources at the bottom of the chart have to be made available or brought online before the linked resources above them. When the cluster is shut down, the resources need to be brought down in the opposite order, from top to bottom. When you use the Application Configuration wizard to create a service group, it establishes these resources.

Prerequisites

- Verify that the binaries of the application to be configured are present on the nodes on which the service group will be configured.
- Verify that the shared drives required by the applications are mounted.
- Before running the wizard, make sure you have the following information ready:
 - Type of applications for which resources are to be configured.
 - Shared storage used by the applications.
 - Registry replication information.
 - Network information.

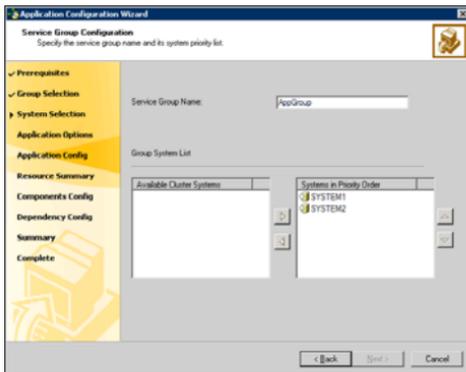
Creating the service group with the application configuration wizard

The steps given in this section may not be the exact steps needed to create a service group for every database application. They are presented to give a typical example of what may be involved in creating a VCS service group for a database application.

To establish a service group for an application database

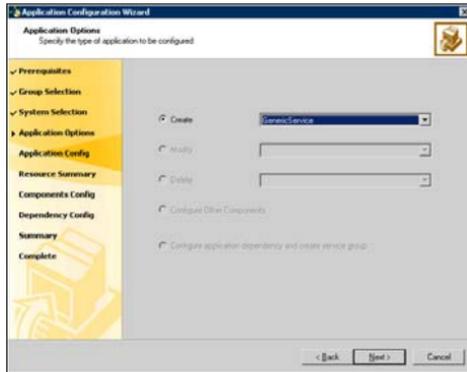
- 1 Select **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Application Configuration Wizard**.
- 2 Read the information in the Welcome panel and click **Next**.
- 3 In the **Wizard Options** panel, click **Create service group** to add a new service group to the cluster. Click **Next** to continue.

4 Complete the following in the Service Group Configuration panel:



- Enter a name for the service group in the **Service Group Name** field. Specify a name that conveys information about the service group you are creating.
 - Select systems on which to configure the service group from the “Available Cluster Systems” list and add them to the “Systems in Priority Order” list. Arrange them in the order that matches their position in the cluster. System priority defines the order in which service groups are failed over to systems. The system at the top of the list has the highest priority.
 - Click **Next** to continue.
- 5 The wizard validates the configuration.

- 6 In the Applications Options panel that appears next, do the following:



- Select **Create** to create a service group for an application.
- Select the agent type used to bring the application online or offline and to monitor its status.

For this example, GenericService would be selected because the database application runs as a Windows service.

If the application does not run as a generic service, you should select the Process agent at this point.

The Service Monitor agent does not control an application. It can be used to monitor a resource that the application resource may depend on. It does not bring the application online or offline. It monitors a service, starts a user-defined script, and interprets the exit code of the script.

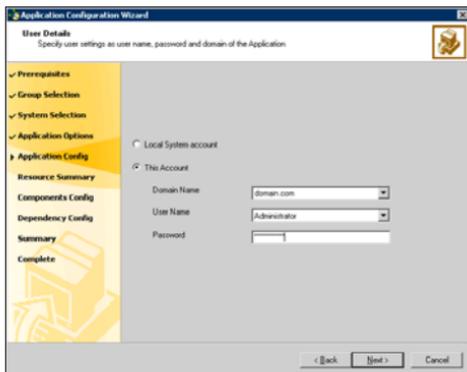
- Click **Next** to continue.

- 7 In the Generic Service Options panel, do the following:



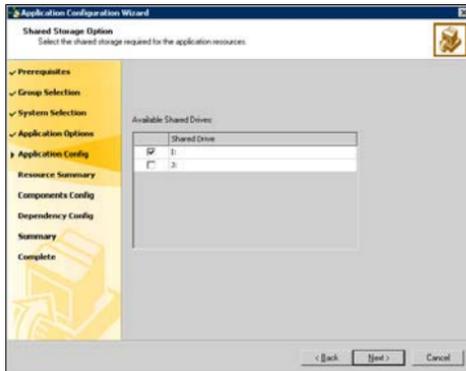
- Select the service name.
Click the icon to the right of the **Service Name** entry box to bring up a list of Windows services and click the desired service. In this example, the service for the database would be selected.
- Provide the start parameters for the service, if applicable.
- In the **Delay After Online** field, specify the number of seconds the agent waits after the service is brought online before starting the monitor routine.
- In the **Delay After Offline** field, specify the number of seconds the agent waits after the service is taken offline before starting the monitor routine.
- Click **Next**.

8 Make the necessary settings in the User Details window:



- Select **This Account** to have the service group run in the cluster.
- Specify the following details about the user.
 - Select the domain in the **Domain Name** list box.
 - Specify the user in the **User Name** list box.
 - Enter the password for the user in the **Password** field.
- Click **Next** to continue.

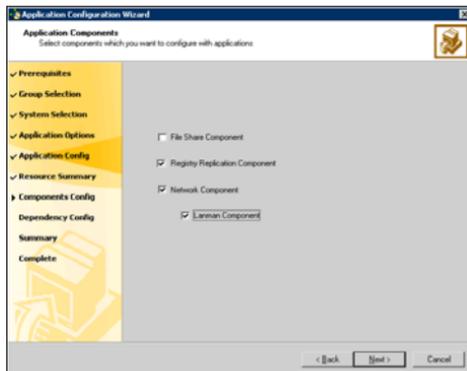
- 9 In the Shared Storage Option window, select the shared storage required for the GenericService resource by clicking the check box adjacent to the shared drive or drives. Click **Next**.



The wizard determines from the storage indicated that SFW cluster disk groups are involved, thus adding the VMDg resource and the MountV resource to the service group.

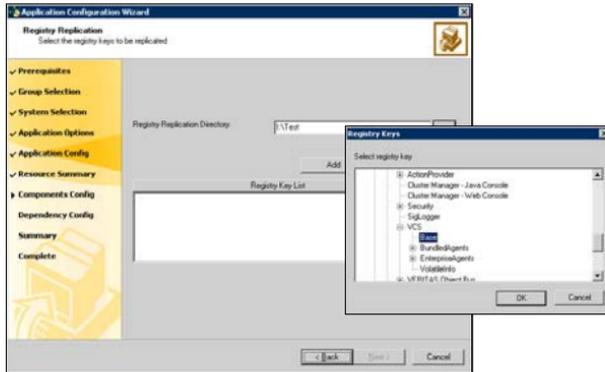
- 10 In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
- 11 The Application Options panel appears. Select the **Configure Other Components** radio button to configure additional resources for the service group and click **Next**.
- 12 In the Application Components window, check the **Registry Replication Component** and **Network Component** check boxes to add the resources for

replicating the registry as well as the NIC, IP, and Lanman resources to the service group. Click **Next**.



- 13 In the Registry Replication window, specify the registry keys to be replicated by doing the following:
 - Specify the directory on the shared disk in which the registry changes are logged.
 - Click **Add**.

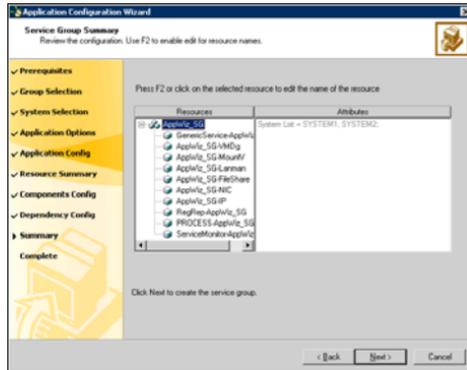
- In the **Registry Keys** panel, select the registry key to be replicated:



- Click **OK**.
 - The selected registry key is added to “Registry Key List” box. Click **Next**.
The RegRep and MountV resources are added to the service group.
- 14 In the Virtual Computer Configuration panel, specify the information related to your network:
- Enter a unique virtual computer name by which the node will be visible to the other nodes. Note that the virtual name must not exceed 15 characters.
- The **Virtual Computer Name** field will not be displayed if you did not select **Lanman Component** in the Application Components panel.
- Enter a unique virtual IP address for the virtual server.
 - Enter the subnet to which the virtual server belongs.
 - For each system in the cluster, select the public network adapter name. To view the adapters associated with a system, click the **Adapter Name** field and click the arrow.
- The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if applicable. Verify that you select the adapters assigned to the public network, not the private network.
- Click **Next**.

The Lanman resource, the virtual IP resource, and the NIC resource are added to the service group.

- 15 In the Application Options panel, click **Configure application dependency and create service group** and click **Next**.
- 16 In the Service Group Summary panel, review your configuration:



- Change the names of resources, if required. The wizard assigns unique names to resources. Click a resource name to edit it. Review your configuration and click **Next**.
 - In the Confirmation dialog box, click **No** to review your settings. Otherwise, click **Yes**. The wizard starts running commands to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion panel appears.
- 17 In the completion dialog box, click **Bring the service group online** check box if you want to bring the service group online on the local system. Click **Finish** to exit the Application Configuration wizard.

Verifying the cluster configuration

To verify the configuration of a cluster, either move the online groups, or shut down an active cluster node.

- Use Veritas Cluster Manager (Java Console) to switch all the service groups from one node to another.

- Simulate a local cluster failover by shutting down an active cluster node.

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in [step 1](#).
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.

The service group you selected is taken offline and brought online on the node that you selected.

Part 2: Setting up the parallel environment on the secondary site

You set up an almost identical configuration on the secondary site. See the following topics:

- [“Installing SFW HA”](#) on page 879
- [“Configuring the VVR security service”](#) on page 879
- [“Configuring the cluster”](#) on page 882
- [“Configuring disk groups and volumes”](#) on page 899

Create an identical disk group and volume setup on the secondary site, as on the primary site. The data volumes on the secondary site should be the same size as the corresponding data volumes on the primary site. The log volume on the secondary site can be a different size, but Symantec recommends that the sizes be the same. The disks, disk groups, and volumes must have the same names. The volumes need to be the same volume types.

- [“Installing the application on cluster nodes”](#) on page 906
- [“Creating VCS service groups”](#) on page 909

Note: Do not bring the service group online if the service group on the primary site is online.

- [“Verifying the cluster configuration”](#) on page 919.

Part 3: Adding the VVR components for replication

This section provides information on configuring the VVR components for replication. Topics include:

- [“VVR components overview”](#) on page 922
- [“Configuring the Replicator Log volumes for VVR”](#) on page 923
- [“Setting up the replicated data sets \(RDS\) for VVR”](#) on page 925
- [“Creating the VVR RVG Service group”](#) on page 936

VVR components overview

The terms Replicated Volume Group (RVG), Replicator Log, and Replicated Data Set (RDS) are used frequently in this section. Here are their definitions:

Replicated Volume Group (RVG)

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, there is a corresponding RVG with a disk group of the same name, and volumes with the same names. The data volumes should be the same size, but Replicator Log volume sizes can differ. Optionally, to add more redundancy, you can have multiple secondary hosts, all with the same corresponding copy of the RVG on the primary host.

An RVG within a disk group is the container for replication, so if you have multiple disk groups, you 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.

Replicated Data Set (RDS)

An RVG on the primary host and any corresponding RVGs on the secondary host or hosts make up a Replicated Data Set (RDS).

Replicator Log

Each RVG must have a Replicator Log associated with it. The Replicator Log volume at the primary site holds a copy of any 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 logs at the two sites must have the same name; however, the sizes of the logs can vary. Symantec recommends

having Replicator Log volumes of the same size at the primary site and the secondary site.

The process described in this section involves setting up an RDS for each SFW disk group on the primary site that will have replicated volumes, and then creating a VVR service group that is linked to the application service group.

Configuring the Replicator Log volumes for VVR

Create the volume for the Replicator Log at each site. The task of creating the logs can also be done during the RDS creation process, but some storage administrators may prefer to do it manually (as is being done here) as a preparatory step to setting up the RDS.

Note: To improve write performance, Symantec recommends that you create the Replicator Log volume on a different disk from the disks used for your application data volumes.

To configure the Replicator Log volumes for VVR

- 1 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.
- 2 Create a volume for the disk group that contains the storage group data:
 - On the System configuration tree, click the disk group where the log volume will be created (*Hostname > Disk Groups > Diskgroupname*).
 - Right-click on the disk group that has the volumes to be replicated, and click **New Volume**.
- 3 On the Welcome panel of the New Volume wizard, click **Next**.
- 4 Select the disks for the volume:
 - Select the group name.
 - Select **Manually select disks**.
 - Click the disk name.
 - Click **Add**.
 - After selecting all the necessary disks, click **Next**.
- 5 Specify the parameters of the volume:
 - Enter the volume name.
 - Enter the size. The size of the Replicator Log volume varies for different environments. To determine the appropriate size for your

environment, refer to the *Veritas Storage Foundation Veritas Volume Replicator, Administrator's Guide*.

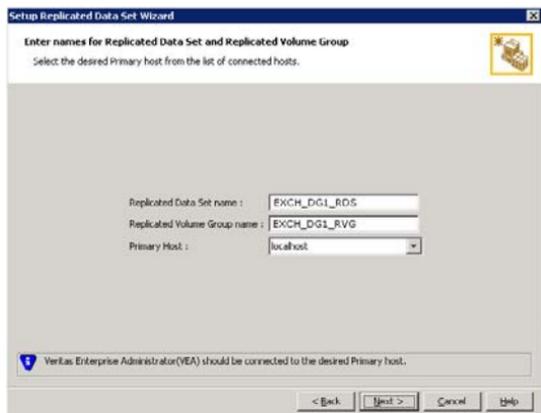
- Select the volume layout.
 - Select the appropriate mirror options.
 - Click **Next**.
- 6 On the Add Drive Letter and Paths dialog box:
- Click **Do not assign a drive letter**.
 - Click **Next**.
- 7 When prompted to format the volume:
- Deselect **Format this volume**.
 - Click **Next**.
- 8 Click **Finish** to create the new volume.
- 9 If necessary, repeat [step 2](#) through [step 8](#) to create Replicator Log volumes for any additional RVGs on the primary site.
- 10 Repeat [step 2](#) through [step 8](#) to create a Replicator Log volume for each RVG on the secondary site.

Setting up the replicated data sets (RDS) for VVR

Configuring VVR involves setting up the Replicated Data Sets on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

To create the Replicated Data Set

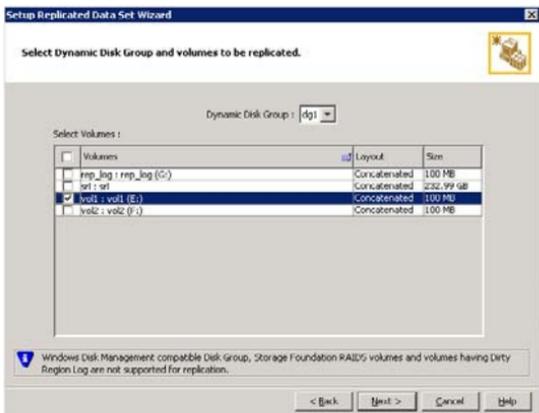
- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported:
Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
From the 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 (APP_REPL_LOG). 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:

The screenshot shows a window titled "Setup Replicated Data Set Wizard" with a sub-header "Edit replication settings". Below the sub-header is the instruction "Edit replication settings or click next." and a small icon of a server rack. The main area contains the following fields:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

At the bottom of the form area is an "Advanced" button. Below the form area is a warning message: "DHCP addresses are not supported by VVR." At the very bottom of the window are navigation buttons: "< Back", "Next >", "Cancel", and "Help".

- 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. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. 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. 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 MISSING.</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.

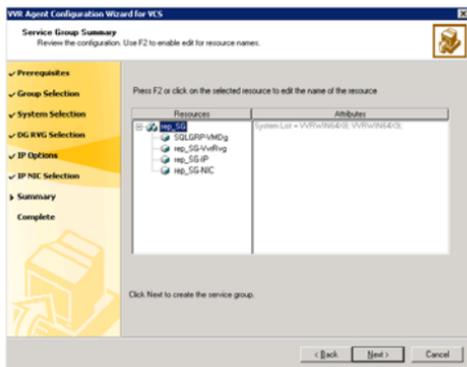
For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- 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 VVR RVG Service group

Run the wizard from the system that contains the application service group.

- Enter the service group name.
- In the Available Cluster Systems box, click the nodes on which to configure the service group, and click the right-arrow icon to move the nodes to the service group's system list. Make sure that the set of nodes selected for the replication service group is the same or a superset of nodes selected for the application's service group. Ensure that the nodes are in the same priority order.



Part 4: Adding GCO components for wide-area recovery

The Global Cluster Option is required to manage global clustering for wide-area disaster recovery. The process of creating a global cluster environment involves the following tasks:

- Connecting standalone clusters by adding a remote cluster to a local cluster.
- Converting the local service group that is common to all the clusters to a global service group.

Use the VCS Java Console or Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to perform global cluster operations; this guide provides procedures only for the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on GCO operations available from the Java and Web Consoles.

Topics in this section include:

- [“Prerequisites for a global cluster environment”](#) on page 937
- [“Linking clusters by adding a remote cluster”](#) on page 938
- [“Converting a local Service group to a global group”](#) on page 939
- [“Additional global cluster administration tasks”](#) on page 942

Prerequisites for a global cluster environment

Creating a global cluster environment requires the following conditions:

- All service groups are properly configured and able to come online.
- The service group that will serve as the global group has the same unique name across all applicable clusters.
- The clusters must use the same version of VCS.
- The clusters must use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.

Linking clusters by adding a remote cluster

The VCS Cluster Manager (Java Console) provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

Note the following uses of the wizard:

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either the cluster on the primary site or the cluster on the secondary site.
- If you are adding a standalone cluster to an existing global cluster environment, run the wizard from a cluster already in the global cluster environment.

The following information is required for the Remote Cluster Configuration Wizard in VCS Cluster Manager:

- The active host name or IP address of each cluster in the global configuration and of the cluster being added to the configuration.
- The user name and password of the administrator for each cluster in the configuration.
- The user name and password of the administrator for the cluster being added to the configuration.

Symantec does not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

To add a remote cluster in Cluster Explorer

- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
- 2 Review the required information for the Remote Cluster Configuration Wizard and then click **Next**.
- 3 In the Wizard Options panel, click **Add Cluster**, then click **Next**.
- 4 In the New Cluster Details panel, enter the details of the new cluster.
If the cluster is not running in secure mode, specify the following:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- If necessary, change the default port number.
- Enter the user name and the password.
- Click **Next**.

If the cluster is running in secure mode, specify the following:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
 - Verify the port number.
 - Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.
 - If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
 - Click **Next**.
- 5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.
 - 6 Verify that the heartbeat connection between clusters is alive. From the command window enter `habb -display`. The state attribute in the output should show **alive**.
If the state is **unknown**, then offline and online the ClusterService group.

Converting a local Service group to a global group

After linking the clusters, use the Global Group Configuration wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator to each cluster in the configuration.

Use the VCS Java Console or Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to bring a global group online, take a global group offline, or switch a global group on a remote cluster; the section below

provides procedures from the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on global cluster operations from the Java Console and Web Consoles.

Note: For remote cluster operations, the user must have the same name and privilege as the user logged on to the local cluster.

To convert a local service group to a global group

- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster, and click **Configure Global Groups**.
or
From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3.
- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify, as follows:
 - Click the name of the service group that will be converted from a local group to a global group, or vice versa.
 - From the Available Clusters box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the Clusters for Service Group box; for global to local cluster conversion, click the left arrow to move the cluster name back to the Available Clusters box. A priority number (starting with 0) indicates the cluster on which the group will attempt to come online. If necessary, double-click the entry in the **Priority** column and enter the new value.
 - Select the policy for cluster failover as follows:

Manual	Prevents a group from automatically failing over to another cluster.
Auto	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
Connected	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.

- Click **Next**.

- 4 Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster, as follows:
- 5 Click **Next**, then click **Finish**.

- | | |
|----------------------------|---|
| Cluster not in secure mode | <ul style="list-style-type: none">■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.■ Verify the port number.■ Enter the user name.■ Enter the password.■ Click OK.■ Repeat these steps for each cluster in the global environment. |
| Cluster in secure mode | <ul style="list-style-type: none">■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.■ Verify the port number.■ Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and domain.■ If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.■ Click OK.■ Repeat these steps for each cluster in the global environment. |

At this point, you must bring the global service group online from Cluster Explorer.

To bring a remote global service group online from Cluster Explorer

- 1 In the Service Groups tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Online**, and click **Remote online**.
- 3 In the Online global group dialog box, specify the following:
 - Click the remote cluster to bring the group online.
 - Click the specific system, or click **Any System**, to bring the group online.

- Click **OK**

Additional global cluster administration tasks

This section provides the following global cluster administration tasks:

- “Taking a remote global Service group offline” on page 942
- “Switching a remote global Service group” on page 942

For further information and procedures relating to global clustering, see the “Global Clustering” section in the *Veritas Cluster Server Administrator’s Guide*.

Taking a remote global Service group offline

To take a remote global service group offline

- 1 On the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Offline**, and click **Remote offline**.
- 3 In the Offline global group dialog box:
 - Click the remote cluster to take the group offline.
 - Click the specific system, or click **All Systems**, to take the group offline and click **OK**.

Switching a remote global Service group

To switch a remote global service group

- 1 On the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 On the Switch global group dialog box:
 - Click the cluster to switch the group.
 - Click the specific system, or click **Any System**, to take the group offline and click **OK**.

Part 5: Maintaining: Normal Operations and recovery procedures

This section provides tasks during normal operations of the DR solutions and also describes the recovery process.

Normal operations: Monitoring the status of the replication

Under normal operating conditions, you can monitor the status of the replication using the following tools:

- The VEA GUI
- The Command Line Interface (CLI)
- Perfmon
- Alerts

For details, refer to the “Monitoring Replication” chapter in the *Veritas Storage Foundation 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. The following are a generic set of tasks that you may need to perform:

- Take the RVG resource offline on both the clusters.
- Transfer the primary role to the host at the secondary site by 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.
- Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- Bring the RVG resource online on the new secondary.
- Bring the application group online on the new primary.

You can now verify that the application functions properly on the new primary with the replicated data. After verifying its functioning, you can revert the roles to what they were originally by repeating the procedure.

Note: Any changes that you make to the data on the new primary will get replicated to the original primary, which is now the secondary.

Disaster recovery procedures

This section provides information on bringing up an application server on the secondary host, in the event of a disaster. It also explains how to migrate the primary role back to the original primary host once it is returned to normal functioning after a disaster.

To bring up the application on the secondary host

- 1 From the left pane in the VEA GUI console on the secondary host, right-click the desired secondary RVG node inside the replication network.
- 2 Select **Takeover** and follow the instructions to perform the takeover operation. You can choose to perform takeover with the following options:
 - Perform **Takeover** *with* the **fast-fallback** option to restore the original primary easily once it becomes available again. When performing **Takeover** with **fast-fallback**, make sure that you do not select the **Synchronize Automatically** option.
 - Perform **Takeover** *without* the **fast-fallback** option. In this case, you will need to perform a complete synchronization of the original primary with the new primary. This may take quite a while, depending on the size of the data volume. Only after the synchronization is complete can you migrate the primary role back to the original primary.

After the takeover, the existing secondary becomes the new primary.

- 3 Assign drive letters to the volumes on the new primary. Make sure that these drive letters are the same as those of the original primary.
- 4 Bring the application group online.

Now you can start using the application on the new primary.

Restoring the primary host

After a disaster, when the original primary becomes available again, you may want to revert the role of the primary back to this host.

To restore the primary host

- 1 Take the RVG resource off-line on both the clusters.

- 2 Depending on whether you performed **Takeover** *with* or *without* the **fast-failback** option, do one of the following:
 - For **Takeover** *with* the **Fast-failback** option:

The original primary, after it has recovered, will be in the **Acting as secondary** state. If the original primary is not in the **Acting as secondary** state, verify whether your network connection has been restored.

To synchronize this original primary and the new primary, use the **Resynchronize Secondaries** option from new primary's context menu.
 - For **Takeover** *without* the **Fast-failback** option:

After performing a takeover without fast-failback, you must convert the original primary to a secondary by 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 RDS's. 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 the **Synchronize Automatically** option.
- 3 After the synchronization is complete, perform a migrate operation to transfer the primary role back to the original primary. Right-click on the primary RVG and select **Migrate** from the menu that appears.
- 4 Make sure that the volumes have retained the same drive letters as they had before the disaster.
- 5 Bring the RVG resource online on the secondary.
- 6 Bring the application group online on the original primary.

Index

A

- ABE 130, 269, 388
- Access-based enumeration 130, 269, 388
- automatic volume growth 848
 - definition 868

C

- campus cluster
 - changing MSCS quorum resource to dynamic quorum resource 714, 757
 - completing setup of application group in MSCS 713, 757
 - configuration 219, 681, 723
 - create SFW cluster disk groups and volumes 700, 740
 - creating VCS service groups 261
 - installing and configuring hardware 688, 730
 - installing application on cluster nodes 257, 711, 754
 - installing Windows and configuring network settings 225
 - making quorum cluster disk group an MSCS resource 709, 750
 - overview 606
 - requirements 212
 - setting up a group for the application 708, 713, 750, 757
 - verifying cluster configuration 320
 - verifying MSCS cluster configuration 714, 757
- campus clustering
 - forceimport attribute of the vmdg resource 224
 - introduction 205
 - MSCS configuration 678, 681, 720, 723
 - MSCS failure scenarios 686, 728
 - prerequisites 678, 720
 - SFW HA configuration 212, 219
 - VCS failure scenarios 221
 - verifying cluster configuration 714, 757
 - Vxclus 688, 730
- cluster

- configure LLT over ethernet 110, 238, 359, 509, 886
- configure LLT over UDP 111, 239, 360, 510, 887
- cluster configuration
 - steps for a new cluster 72
 - steps for an existing cluster 72
- cluster node
 - installing application 257, 635, 668, 711, 754
- clustering concepts
 - ownership of quorum 687, 729
 - quorum 687, 729
- clusters
 - assigning user privileges 528
 - configuring (RDC) 355
 - switching online nodes 469
 - verifying configuration 181, 919
- cmdlets 582
- configure
 - LLT over ethernet 110, 238, 359, 509, 886
 - LLT over UDP using VCW 111, 239, 360, 510, 887
- configure cluster
 - ethernet 886
 - UDP 887
- configure cluster
 - ethernet 110, 238, 359, 509
 - UDP 111, 239, 360, 510
- configuring
 - cluster (RDC) 355
 - network and storage 343
- configuring network and storage 343
- creating new volumes 701, 742

D

- detail monitoring
 - configuring 200
 - disabling 202
 - enabling 201
- disaster recovery
 - about 474, 475

- adding GCO components for wide-area recovery 937
- adding VVR components for replication 784, 823, 922
- changing MSCS quorum resource to dynamic quorum resource 778
- completing setup of application group in MSCS 776
- components of VVR that enable disaster recovery 477
- configuring replicator log volumes for VVR 785, 824, 923
- creating VVR RVG service group 936
- defined 474
- disk space requirements 485
- establishing cluster under MSCS (primary site) 769, 814
- illustrated 476
- installing and configuring hardware 767, 814
- installing SFW (primary site) 769, 816
- normal operations and recovery procedures 563, 568, 800, 840, 943
- overview 473, 476
- setting up cluster on primary site 763, 810, 879
- setting up cluster on secondary site 781, 821, 921
- SFW-Microsoft failover cluster-VVR configuration 805
- SFW-MSCS-VVR configuration 759
- solution 475
- disaster recovery (DR)
 - configuring GCO with DR wizard 546
- Disaster Recovery Procedures 572, 801, 841, 944
- disk groups
 - deporting 104, 377
 - importing 104, 377
 - overview 97
- disk groups and volumes
 - configuring 97
 - creating SFW cluster disk groups 249, 700, 740
 - managing 104, 377
- disk space requirements 485, 614, 647, 680, 722, 765, 812
- DMP DSMs 172
- DR
 - multiple sites 567
- DR wizard
 - configuring replication and GCO 546
- driver signing options

- resetting 96, 233, 352, 504
- dynamic multi-pathing
 - about 68
 - adding to a clustering configuration 67
 - configuration 71
 - configuration tasks 69
 - existing cluster configuration 72
 - more on DMP paths 862
 - new cluster configuration 72
 - prerequisites 69
 - server consolidation 862
- dynamic quorum
 - implementation 637, 671

E

- EMC requirements 523

F

- Failover clustering
 - campus clustering configuration 643
 - local clustering configuration 643
- FastResync (FR) 54
- FileShare
 - access-based enumeration 130, 269, 388
- Fire Drill Wizard
 - actions 589
 - changing a fire drill configuration 597
 - deleting the configuration 600
 - overview 578
 - preparing the configuration 589
 - prerequisites for a fire drill 586, 587, 588
 - restoring the prepared configuration 599
 - running a fire drill 595
- FlashSnap 53
 - quick recovery 53
 - reporting and analysis 59
 - tips and references 65
- forceimport attribute of vmdg resource 224

G

- Global Cluster Option (GCO)
 - configuring with the DR wizard 546
 - secure configuration 565
- global clusters
 - adding 938
 - prerequisites 937

H

- hardware, configuring 343
- high availability
 - about 75, 605
 - local clustering and high availability 75
 - overview 606
 - solution 75, 605
- Hitachi TrueCopy
 - requirements for DR wizard 526

I

- IIS configuration, synchronizing 122, 147, 259, 286, 380, 405
- installing
 - applications, secondary site tasks 435
 - verifying the installation 434
- installing and configuring hardware 224, 688, 730

L

- LLT over ethernet
 - configuring using VCW 110, 238, 359, 509, 886
- LLT over UDP
 - configuring using VCW 111, 239, 360, 510, 887
- local clustering
 - high availability 75
 - MSCS 609, 643
 - overview 75

M

- Microsoft volume shadow copy service (VSS) 55
- modes of replication
 - asynchronous 478
 - synchronous 477
 - synchronous override 478
- MSCS
 - campus clustering configuration 609
 - changing to a dynamic quorum 714, 757
 - completing setup of application group 636, 670
 - configuration 615, 649
 - configuring the network and storage 616, 650
 - creating dynamic cluster disk groups 628, 661
 - creating dynamic volumes 630, 663
 - disaster recovery procedures 801, 841
 - establishing cluster 690, 732
 - establishing cluster under MSCS 690, 732
 - installing application on cluster nodes 635,

668

- installing SFW 618, 653
- local clustering configuration 609
- prerequisites 612, 646
- setting up a group for application 633, 667, 708, 713, 750, 757
- setting up a group for the application 633, 667, 773
- setup of application group 636, 670, 713, 757
- verifying cluster configuration 640, 673
- multiple disk group best practices
 - disk group clusters 43
 - disk group structure 42
 - quorum device configuration 607
- multiple DR sites 567

N

- network and storage
 - configuration 616, 650
- network, configuring 343
- new volumes
 - creating 701, 742

O

- off-host backup, defined 58
- online storage migration 848

P

- performance monitoring
 - data-transfer intensive applications 46
 - request intensive applications 45
- permissions requirements 85, 217, 341, 491
- post-fire drill scripts 582
- print share groups
 - modifying using wizard 194

Q

- quick recovery
 - about 51
 - components 53
 - definition 51
 - example 62
 - Oracle database example 62
 - overview 56
 - recover database using split-mirror snapshot and database logs 63
 - solution 51

understanding components of quick recovery 53

R

RAID best practices
 hardware RAID 48
 mirroring 40, 41
 RAID-5 44
 read performance and failure tolerance 43
 striping across hardware 44
 striping and mirroring 41
 RAID Configurations Using Logical Volume Management 40, 43
 replicated data clusters
 setting up 328
 replicated data clusters (RDC)
 setting up a secondary site 435
 replicated data set (RDS) 922
 replicated volume group (RVG) 922, 936
 replication
 adding VVR components 784, 823, 922
 asynchronous 478
 configuring with the DR wizard 546
 general definition 477
 replicator log 922
 RVG snapshot 479
 setting up a Replicated Data Set (RDS) 441
 synchronous 477
 synchronous override 478
 write order fidelity 478
 requirements
 disk space 614, 647, 680, 722, 765, 812
 permissions 85, 217, 341, 491
 requirements, additional for SFW HA 85, 217, 341, 491
 requirements, network 84, 216, 340, 490
 requirements, system 84, 216, 340, 490
 resetting
 driver signing options 96, 233, 352, 504

S

SAN
 setting up storage 853
 SFW features 854, 868
 scripts 582
 secure clusters
 assigning user privileges 528
 secure GCO, establishing 565

Security Services
 configuring 113, 241, 362, 512, 889
 server consolidation 845
 about 847
 adding storage array 859
 advantages 848
 advantages of using SFW 848
 configuration 852, 856, 863
 configuration 1 – many to one 853
 configuration 2 – many to two
 adding clustering and DMP 860
 customer success story 869
 definition 847
 general configuration 852
 many-to-one configuration 853
 many-to-two configuration with MSCS and DMP 860
 migrating the data to the large server 858
 overview 847, 850
 performance monitoring 849, 868, 869
 preparing to consolidate 857
 process to implement 850
 SFW features 868

Server Core
 configure file shares 125, 264, 383
 configuring IIS sites 146, 285, 404
 installing IIS 147, 286, 405

service groups
 administering global groups 939
 creating 909
 dependencies 493, 574
 setting bandwidth
 using RDS wizard 551

setting up cluster
 primary site 763, 810, 879
 secondary site 781, 821, 921

SFW
 best practices 39
 high availability configuration 75
 implementing dynamic MSCS quorum resource 637, 671
 installing 618, 653
 SFW-specific solutions 23, 203
 typical high availability configuration 75

SFW HA
 additional requirements 85, 217, 341, 491
 best practices 86, 218, 342, 492
 creating cluster disk groups and volumes 249
 creating dynamic volumes 101, 373, 903

- disaster recovery procedures 572, 944
- installing 228
- installing and configuring hardware 224, 493
- installing application (primary site) 906
- network requirements 84, 216, 340, 490
- service group example 909
- system requirements 84, 216, 340, 490
- verifying cluster configuration 320
- snapshot
 - commands 53
 - other applications for point-in-time snapshots 58
- Solutions Configuration Center
 - context sensitivity 31
 - overview 29
 - running wizards remotely 34
 - starting 30
 - wizard descriptions 34
- SQL Server Virtual Device Interface (VDI)
 - quiescing the database 52
- SRDF requirements 523
- storage capacity best practices
 - allocation planning 47
 - failure-tolerant volume recovery 48
 - location of data objects 40
 - manage unallocated space 48
- storage configuration 343
- switching online nodes 469

T

- tips
 - FlashSnap 65

U

- user privileges
 - assigning 528

V

- VCS
 - campus clustering configuration 209
 - configuring the cluster for RDC 355
 - switching online nodes 469
 - VVR configuration 481
- VCS service groups
 - creating 261
- verifying
 - cluster configuration for HA 181, 919

volumes

- creating 704, 745
- creating on primary 253, 704, 745
- mounting 104, 377
- overview 97
- unmounting 104, 377

VSS 55

VVR

- components that enable disaster recovery 477
- configuring replication with DR wizard 546
- setting up RDS (VCS) 441
- setting up Replicated Data Sets (RDS) 925
- SFW HA-VCS configuration 481, 873
- Vxclus utility 688, 730

W

- wide-area recovery
 - adding GCO components 937
- Windows
 - network settings 225

