Storage Foundation and High Availability Solutions, Solutions Guide

Windows

7.1

April 2016

VERITAS™

Storage Foundation and High Availability Solutions, Solutions Guide

Document version: 7.1 Rev 0

Last updated: 2016-04-25

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

.

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

doc_feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1 Introduction 19

Chapter 1 Introducing Storage Foundation and High Availability Solutions 20

About the solutions guides 20
About Quick Recovery 21
About high availability 21
About campus clusters 21
About disaster recovery 21
About Microsoft clustering solutions 22
How this guide is organized 22

Chapter 2 Using the Solutions Configuration Center 23

About the Solutions Configuration Center 23
Starting the Solutions Configuration Center 24
Options in the Solutions Configuration Center 24
About launching wizards from the Solutions Configuration Center 25
Remote and local access to Solutions wizards 26
Solutions wizards and logs 28
Workflows in the Solutions Configuration Center 28

Chapter 3 SFW best practices for storage 30

Best practices for storage availability 30
Adding software mirrors for critical data 30
Locating data objects for optimum recovery 31
Managing three-way software mirrors for reliability 31
Software striping and mirroring on top of hardware RAID for high availability 31
Best practices configuring SFW disk groups for availability 32
Configuring disk groups for separate storage capacity pools or common pools 32
Allocating disk groups for availability in clusters 33
Best practices for storage performance 33

Host-based mirroring for increased read performance and failure tolerance 33
Software striping across hardware for increased performance 34
Using software RAID 5 for read-mostly data 34
Best practices for I/O performance tuning 34
Striping for I/O-request-intensive applications 35
Striping for data-transfer-intensive applications 36
Best practices for storage capacity management 36
Managing storage allocation for flexibility 37
Aggregating hardware RAID for very large volumes 37
Managing unallocated space for free space savings 37
Reserving spares for failure-tolerant volume recovery 38

Section 2 Quick Recovery 39

Chapter 4 Quick Recovery overview 40

About the Quick Recovery solution 40
Need for implementing the SFW Quick Recovery solution 41
Understanding the underlying components of SFW's Quick Recovery process 42
Overview of the Quick Recovery process 44
Creating initial snapshots 44
Refreshing a snapshot 44
Recovering a database 45
Other applications for point-in-time snapshots 45
Off-host backups 46
Reporting and analysis 47
Application testing and training 48

Chapter 5 Quick Recovery example 50

Example of quick recovery of an Oracle database 50
Creating a split-mirror snapshot of the database 51
Recovering the database using the split-mirror snapshot and database logs 52
Tips and references about FlashSnap 54

Section 3 High Availability 55

Chapter 6 High availability: Overview 56

About high availability 56
About clusters 56
How VCS monitors storage components 57
Shared storage—if you use NetApp filers 57
Shared storage—if you use SFW to manage cluster dynamic disk groups 58
Shared storage—if you use Windows LDM to manage shared disks 58
Non-shared storage—if you use SFW to manage dynamic disk groups 59
Non-shared storage—if you use Windows LDM to manage local disks 59
Non-shared storage—if you use VMware storage 60

Chapter 7 Deploying InfoScale Enterprise for high availability: New installation 61

About the high availability solution 62
Tasks for a new high availability (HA) installation—additional applications 62
Reviewing the InfoScale installation requirements 64
Notes and recommendations for cluster and application configuration 330
IPv6 support 333
Reviewing the configuration 67
Configuring the storage hardware and network 335
About installing the Veritas InfoScale products 503
Configuring disk groups and volumes 70
Planning disk groups and volumes 71
Creating a dynamic disk group 508
Creating dynamic volumes 475
About managing disk groups and volumes 337
Importing a disk group and mounting a volume 337
Unmounting a volume and deporting a disk group 338
Configuring the cluster using the Cluster Configuration Wizard 345
Configuring notification 354
About modifying the cluster configuration 91
Adding nodes to a cluster 271
Removing nodes from a cluster 96

Reconfiguring a cluster 98
Configuring the ClusterService group 102
Deleting a cluster configuration 106
About installing and configuring the application or server role 108
About configuring a File Share server role 108
About configuring a Print Share server role 108
About installing and configuring the IIS application 109
About installing additional applications 109
Configuring the service group 110
About configuring file shares 110
About configuring print shares 124
About configuring IIS sites 132
About configuring applications using the Application Configuration Wizard 143
About configuring the Oracle service group using the wizard 156
Enabling fast failover for disk groups (optional) 162
Configuring the service group in a non-shared storage environment 163
Verifying the cluster configuration 269
Possible tasks after completing the configuration 166
Modifying the application service groups 166
Modifying a file share service group using the wizard 167
Modifying a print share service group using the wizard 168
Modifying an IIS service group using the wizard 169
Modifying an application service group 170
About modifying an Oracle service group 172

Chapter 8 Adding DMP to a clustering configuration 177

About Dynamic Multi-Pathing 177
Overview of configuration tasks for adding DMP DSMs 178
Reviewing the prerequisites 178
Reviewing the configuration 179
Setting up DMP in a new cluster configuration 180
Setting up DMP in an existing cluster configuration 180

Section 4 Campus Clustering 181

Chapter 9 Introduction to campus clustering 182

About Campus Clusters 182
Sample campus cluster configuration 183
Differences between campus clusters and local clusters 184

Chapter 10 Deploying InfoScale Enterprise for campus cluster 185

About the Campus Cluster solution 185
Notes and recommendations for cluster and application configuration 330
IPv6 support 333
Campus cluster requirements 190
Reviewing the configuration 190
Overview of campus clustering with VCS 192
Reinstating faulted hardware 193
Setting the ForceImport attribute 194
Installing and configuring the hardware 195
Configuring the storage hardware and network 335
About installing the Veritas InfoScale products 503
Configuring the cluster using the Cluster Configuration Wizard 345
Configuring notification 354
Creating disk groups and volumes 210
About cluster disk groups and volumes 210
Example disk group and volume configuration in campus cluster 211
Considerations when creating disks and volumes for campus clusters 506
Viewing the available disk storage 213
Creating a dynamic disk group 508
Adding disks to campus cluster sites 509
Creating volumes for campus clusters 510
Installing the application on cluster nodes 220
About configuring a File Share server role in a campus cluster 221
About configuring a Print Share server role in a campus cluster 221
About installing and configuring the IIS application in a campus cluster 222
About installing additional applications in a campus cluster 222
Deporting and importing a disk group in a campus cluster 223
Configuring service groups 224
Verifying the cluster configuration 269

Section 5 Replicated Data Clusters 227

Chapter 11 Introduction to Replicated Data Clusters 228

About Replicated Data Clusters 228
How VCS Replicated Data Clusters work 229
Setting up a Replicated Data Cluster configuration 231
Setting up replication 231
Configuring the service groups 232
Migrating the service group 233

Chapter 12 Deploying Replicated Data Clusters: New application installation 235

Tasks for a new replicated data cluster installation—additional applications 236
Notes and recommendations for cluster and application configuration 330
IPv6 support 333
Sample configuration 241
Configuring the storage hardware and network 335
About installing the Veritas InfoScale products 503
Setting up security for Volume Replicator 531
Configuring the cluster using the Cluster Configuration Wizard 345
Configuring notification 354
Configuring disk groups and volumes 258
Planning disk groups and volumes 258
Creating a dynamic disk group 508
Creating dynamic volumes 475
About managing disk groups and volumes 337
Installing and configuring the application or server role 266
Configuring a File Share server role 266
Configuring a Print Share server role 266
Installing and configuring the IIS application 267
Installing additional applications 267
Configuring the service group 268
Creating the primary system zone for the application service group 268
Verifying the cluster configuration 269
Creating a parallel environment in the secondary zone 270
Adding nodes to a cluster 271
Creating the Replicated Data Sets with the wizard 539
Configuring a RVG service group for replication 288

Creating the RVG service group 288
Configuring the resources in the RVG service group for RDC replication 289
Configuring the RVG Primary resources 297
Configuring the primary system zone for the RVG service group 299
Setting a dependency between the service groups 299
Adding the nodes from the secondary zone to the RDC 300
Adding the nodes from the secondary zone to the RVG service group 300
Configuring secondary zone nodes in the RVG service group 302
Configuring the RVG service group NIC resource for fail over (VMNSDg only) 302
Configuring the RVG service group IP resource for failover 304
Configuring the RVG service group VMNSDg resources for fail over 306
Adding the nodes from the secondary zone to the application service group 307
Configuring the zones in the application service group 309
Configuring the application service group IP resource for fail over (VMNSDg only) 310
Configuring the application service group NIC resource for fail over (VMNSDg only) 311
Verifying the RDC configuration 312
Bringing the service group online 312
Switching online nodes 312
Additional instructions for GCO disaster recovery 313

Section 6 Disaster Recovery 315

Chapter 13 Disaster recovery: Overview 316

About a disaster recovery solution 316
Need for implementing a disaster recovery solution 318
Overview of the recovery process 318
Components of Volume Replicator that enable disaster recovery 319
Understanding replication 319
Modes of replication 320
Features of Volume Replicator that help in disaster recovery 321

Chapter 14 Deploying disaster recovery: New application installation 323

Tasks for a new disaster recovery installation—additional applications 324
Tasks for setting up DR in a non-shared storage environment 327
Notes and recommendations for cluster and application configuration 330
IPv6 support 333
Reviewing the configuration 333
Supported disaster recovery configurations for service group dependencies 334
Configuring the storage hardware and network 335
About managing disk groups and volumes 337
Importing a disk group and mounting a volume 337
Unmounting a volume and deporting a disk group 338
Setting up the secondary site: Configuring the InfoScale product and setting up a cluster 338
About installing the Veritas InfoScale products 503
Configuring the cluster using the Cluster Configuration Wizard 345
Configuring notification 354
Verifying that your application or server role is configured for HA at the primary site 357
Setting up your replication environment 357
Setting up security for Volume Replicator 531
Requirements for EMC SRDF array-based hardware replication 360
Requirements for Hitachi TrueCopy array-based hardware replication 362
Assigning user privileges (secure clusters only) 364
About configuring disaster recovery with the DR wizard 365
Configuring disaster recovery with the DR wizard 367
Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option) 370
Creating temporary storage on the secondary site using the DR wizard (array-based replication) 374
Installing and configuring the application or server role (secondary site) 379
Installing the FileShare application 379
Installing the PrintShare application 379
Installing the IIS application 380
Installing additional applications 380

Cloning the service group configuration from the primary site to the secondary site 381
Configuring the application service group in a non-shared storage environment 385
Configuring replication and global clustering 385
Configuring Volume Replicator replication and global clustering 386
Configuring EMC SRDF replication and global clustering 394
Configuring Hitachi TrueCopy replication and global clustering 397
Configuring global clustering only 401
Creating the replicated data sets (RDS) for Volume Replicator replication 404
Creating the Volume Replicator RVG service group for replication 404
Configuring the global cluster option for wide-area failover 405
Linking clusters: Adding a remote cluster to a local cluster 406
Converting a local service group to a global service group 407
Bringing a global service group online 409
Verifying the disaster recovery configuration 410
Establishing secure communication within the global cluster (optional) 412
Adding multiple DR sites (optional) 414
Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment 414
Preparing the new node 415
Preparing the existing DR environment 415
Modifying the replication and application service groups 416
Reversing replication direction 417
Maintaining: Normal operations and recovery procedures (Volume Replicator environment) 417
Normal operations: Monitoring the status of the replication 417
Performing planned migration 418
Disaster recovery procedures 418
Recovery procedures for service group dependencies 420

Chapter 15 Testing fault readiness by running a fire drill 423

About disaster recovery fire drills 423
About the Fire Drill Wizard 424
About Fire Drill Wizard general operations 424
About Fire Drill Wizard operations in a Volume Replicator environment 425

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment 428
About post-fire drill scripts 430
Tasks for configuring and running fire drills 431
Prerequisites for a fire drill 433
Prerequisites for a fire drill in a Volume Replicator environment 434
Prerequisites for a fire drill in a Hitachi TrueCopy environment 435
Prerequisites for a fire drill in an EMC SRDF environment 435
Preparing the fire drill configuration 436
System Selection panel details 438
Service Group Selection panel details 439
Secondary System Selection panel details 439
Fire Drill Service Group Settings panel details 439
Disk Selection panel details 439
Hitachi TrueCopy Path Information panel details 440
HTCSnap Resource Configuration panel details 441
SRDFSnap Resource Configuration panel details 441
Fire Drill Preparation panel details 442
Running a fire drill 443
Re-creating a fire drill configuration that has changed 445
Restoring the fire drill system to a prepared state 448
Deleting the fire drill configuration 449
Fire Drill Deletion panel details 450
Considerations for switching over fire drill service groups 451

Section 7 Microsoft Clustering Solutions 452

Chapter 16 Microsoft clustering solutions overview 453

About Microsoft clustering with high availability 453
About Microsoft clustering with Volume Replicator 454
About Microsoft clustering with campus clustering 454
About the SFW-Microsoft clustering-Volume Replicator configuration 454
Configuring the quorum device for high availability 455

Chapter 17 Deploying SFW with Microsoft failover clustering 457

Tasks for deploying InfoScale Storage with Microsoft failover clustering 458
Reviewing the configuration 459
Configuring the storage hardware and network 498
Establishing a Microsoft failover cluster 529
Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering 502
Pre-installation task: moving the online groups 502
About installing the Veritas InfoScale products 503
Installing the server components using the installation wizard 465
Installing the client components 471
Post-installation task: moving the online groups 503
Creating SFW disk groups and volumes 472
Planning disk groups and volumes 473
Creating a dynamic disk group 508
Creating dynamic volumes 475
Creating a group for the application in the failover cluster 478
Installing the application on cluster nodes 479
Completing the setup of the application group in the failover cluster 480
Implementing a dynamic quorum resource 481
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource 482
Adding a VMDg resource for the quorum 515
Changing the quorum resource to a dynamic mirrored quorum resource 516
Verifying the cluster configuration 520
Configuring InfoScale Storage in an existing Microsoft Failover Cluster 485

Chapter 18 Deploying SFW with Microsoft failover clustering in a campus cluster 487

Tasks for deploying InfoScale Storage with Microsoft failover clustering in a campus cluster 488
Reviewing the configuration 489
Overview of campus clustering with Microsoft clustering 491
Campus cluster failure with Microsoft clustering scenarios 492
Microsoft clustering quorum and quorum arbitration 496

Configuring the storage hardware and network 498
Establishing a Microsoft failover cluster 529
Connecting the two nodes 501
Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering 502
Pre-installation task: moving the online groups 502
About installing the Veritas InfoScale products 503
Post-installation task: moving the online groups 503
Creating disk groups and volumes 504
Example disk group and volume configuration in campus cluster 505
Considerations when creating disks and volumes for campus clusters 506
Viewing the available disk storage 507
Creating a dynamic disk group 508
Adding disks to campus cluster sites 509
Creating volumes for campus clusters 510
Implementing a dynamic quorum resource 514
Creating a dynamic cluster disk group and a mirrored volume for the quorum resource 515
Adding a VMDg resource for the quorum 515
Changing the quorum resource to a dynamic mirrored quorum resource 516
Setting up a group for the application in the failover cluster 517
Installing the application on the cluster nodes 518
Pointers for installing the application on the first node 518
Pointers for installing the application on the second node 518
Completing the setup of the application group in the cluster 519
Verifying the cluster configuration 520

Chapter 19 Deploying SFW and VVR with Microsoft failover clustering 522

Tasks for deploying InfoScale Storage and Volume Replicator with Microsoft failover clustering 522
Part 1: Setting up the cluster on the primary site 526
Reviewing the prerequisites and the configuration 526
Installing and configuring the hardware 529
Installing Windows and configuring network settings 529
Establishing a Microsoft failover cluster 529
Installing InfoScale Storage (primary site) 531
Setting up security for Volume Replicator 531
Creating SFW disk groups and volumes 534

Completing the primary site configuration 535
Part 2: Setting up the cluster on the secondary site 535
Repeating cluster configuration steps for the secondary site 536
Part 3: Adding the Volume Replicator components for replication 537
Volume Replicator components overview 537
Configuring the Replicator Log volumes for Volume Replicator 538
Creating the Replicated Data Sets with the wizard 539
Creating resources for Volume Replicator 552
Creating an RVG resource and setting the dependencies 552
Part 4: Maintaining normal operations and recovery procedures 555
Normal operations: Monitoring the status of the replication 556
Performing planned migration 556
Disaster recovery procedures 556

Section 8 Server Consolidation 560

Chapter 20 Server consolidation overview 561

Server consolidation definition 561
Need for implementing server consolidation 561
Advantages of using SFW with server consolidation 562
Overview of the server consolidation process 563

Chapter 21 Server consolidation configurations 565

Typical server consolidation configuration 565
Proof of concept 566
Server consolidation configuration 1—many to one 567
About this configuration 567
Preparing to consolidate 569
Migrating the data to the large server 570
Adding the storage array 571
Completing the consolidation process 572
Server consolidation configuration 2—many to two: Adding clustering and DMP 572
About this configuration 573
Adding the new hardware 575
Establishing the Microsoft failover cluster 576
Adding SFW support to the cluster 576
Setting up Microsoft failover cluster groups for the applications 577
Installing applications on the second computer 578

Completing the setup of the application group in the Microsoft cluster 578
Changing the quorum resource to the dynamic quorum resource 578
Verifying the cluster configuration 578
Enabling DMP 578
SFW features that support server consolidation 579
Automatic volume growth 579
Features that support storage in a SAN 580
Performance monitoring 580

Section 1

Introduction

- Chapter 1. Introducing Storage Foundation and High Availability Solutions
- Chapter 2. Using the Solutions Configuration Center
- Chapter 3. SFW best practices for storage

Chapter 1

Introducing Storage Foundation and High Availability Solutions

This chapter includes the following topics:

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

About the solutions guides

The *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 Server 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.

Storage Foundation and High Availability Solutions (SFW HA) includes Storage Foundation and 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.

Storage Foundation and High Availability Solutions (SFW HA) provides the capability for implementing disaster recovery.

About Microsoft clustering solutions

Microsoft clustering may be used with Storage Foundation to provide high availability for any application or server role.

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

How this guide is organized

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

See "About the Solutions Configuration Center" on page 23.

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.

Chapter 2

Using the Solutions Configuration Center

This chapter includes the following topics:

- About the Solutions Configuration Center
- Starting the Solutions Configuration Center
- Options in the Solutions Configuration Center
- About launching wizards from the Solutions Configuration Center
- Remote and local access to Solutions wizards
- Solutions wizards and logs
- Workflows in the Solutions Configuration Center

About the Solutions Configuration Center

The Storage Foundation and High Availability Solutions Configuration Center guides you through setting up your Storage Foundation (SFW) or Storage Foundation and High Availability Solutions (SFW HA) environment.

The Solutions Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2007 and 2010
- Microsoft SQL Server 2008, 2008 R2, and 2012
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server 2010 and 2013 (high availability, disaster recovery, and Quick Recovery solutions)
- Additional applications

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

- High availability at a single site for a new installation
- High availability at a single site for an existing server
- Campus cluster disaster recovery, including the following:
  - Campus cluster using 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
- Fire drill to test the fault readiness of a disaster recovery environment

Starting the Solutions Configuration Center

Depending on the operating system, you can start the Solutions Configuration Center from the **All Programs** menu, the **Run** menu, or from the **Apps** menu.

To start the Solutions Configuration Center

◆ Click **Start > All Programs > Veritas > Veritas Storage Foundation > Solutions Configuration Center**.

or

Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**.

or

Click **Start > Run**, type **scc**, and press Enter.

or

Navigate to the Apps menu and then click **SCC**.

Options in the Solutions 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 steps that are shown when you click on a solution are customized for that application.

The following figure shows the solutions available when you click Solutions for Additional Applications.

Figure 2-1 Solutions Configuration Center for additional applications

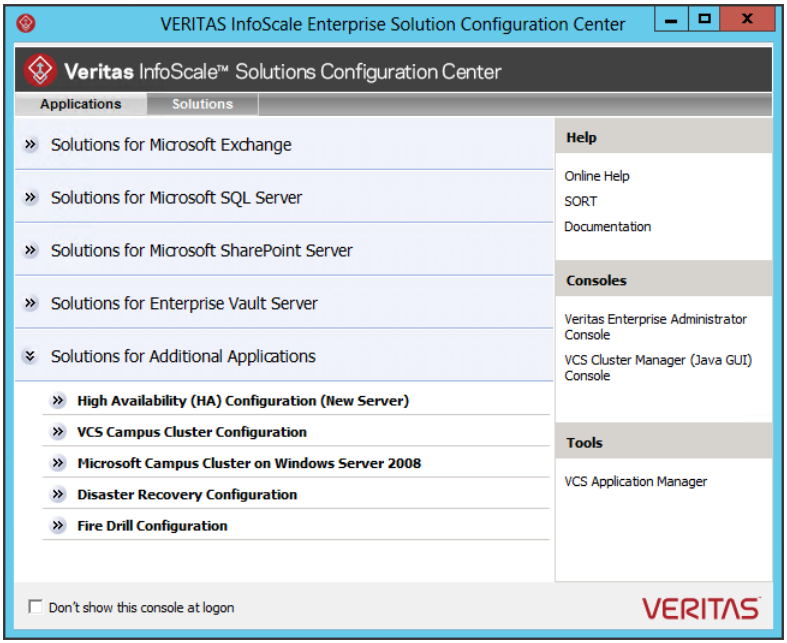

About launching wizards from the Solutions Configuration Center

The Solutions Configuration Center provides two ways to access wizards:

Applications — Lists solutions by application. Provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.

Solutions	(For advanced users) Lists wizards by solution, without additional instructions, as follows: ■ High Availability Configuration Wizards ■ Disaster Recovery Configuration Wizards ■ Quick Recovery Configuration Wizards ■ Fire Drill Configuration Wizards You can go directly to a particular wizard. **Note:** Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems. The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems. Refer to the following Microsoft knowledge database article for more details: http://technet.microsoft.com/en-us/library/dd184075.aspx

Remote and local access to Solutions wizards

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

Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster Can also configure: ■ Volume Replicator (Volume Replicator) replication ■ VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication **Note:** Requires first configuring high availability on the primary site. To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.

Fire Drill Wizard	Sets up a fire drill to test disaster recovery **Note:** Requires first configuring high availability on the primary site. To configure IPv6 settings, launch the wizard from a system on which the IPv6 stack is installed.
Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots
VCS Configuration Wizard	Sets up the VCS cluster
Volume Replicator Security Service Configuration Wizard	Configures the Volume Replicator security service

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
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group
SFW Configuration Utility for Hyper-V Live Migration Support	Configures SFW for Microsoft Hyper-V Live Migration support on the selected systems

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 **Note:** The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems using the Print Share Configuration wizard.

IIS Configuration Wizard	Configures IIS 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

Solutions wizards and 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 folder:

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

Workflows in the Solutions 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 Solutions 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.

The following figure shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

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

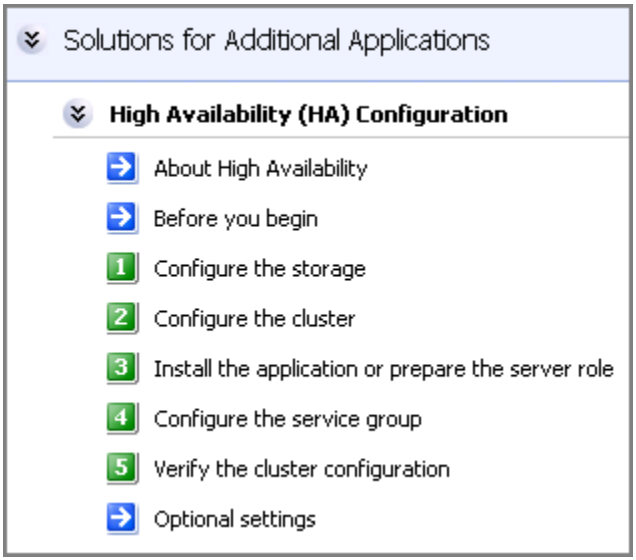

Chapter 3

SFW best practices for storage

This chapter includes the following topics:

- Best practices for storage availability
- Best practices configuring SFW disk groups for availability
- Best practices for storage performance
- Best practices for I/O performance tuning
- Best practices for storage capacity management

Best practices for storage availability

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

- See "Adding software mirrors for critical data" on page 30.
- See "Locating data objects for optimum recovery" on page 31.
- See "Managing three-way software mirrors for reliability" on page 31.
- See "Software striping and mirroring on top of hardware RAID for high availability" on page 31.

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, database recovery is impossible, because both data and logs are inaccessible.

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.

Note that if a mirror is regularly broken off from a three-way mirrored volume, susceptibility to failure is greater during the interval between the third-mirror break-off 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

Storage Foundation (SFW) 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:

- See "Configuring disk groups for separate storage capacity pools or common pools" on page 32.
- See "Allocating disk groups for availability in clusters" on page 33.

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:

- See "Host-based mirroring for increased read performance and failure tolerance" on page 33.
- See "Software striping across hardware for increased performance" on page 34.
- See "Using software RAID 5 for read-mostly data" on page 34.

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 35.

- 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 36.

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:

- See "Managing storage allocation for flexibility" on page 37.
- See "Aggregating hardware RAID for very large volumes" on page 37.

- See "Managing unallocated space for free space savings" on page 37.
- See "Reserving spares for failure-tolerant volume recovery" on page 38.

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.

Section 2

Quick Recovery

- Chapter 4. Quick Recovery overview
- Chapter 5. Quick Recovery example

Chapter 4

Quick Recovery overview

This chapter includes the following topics:

- About the Quick Recovery solution
- Need for implementing the SFW Quick Recovery solution
- Understanding the underlying components of SFW's Quick Recovery process
- Overview of the Quick Recovery process
- Other applications for point-in-time snapshots

About the Quick Recovery solution

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 Server databases is supported.

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

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 *Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft Exchange*.

For detailed information about implementing a Quick Recovery solution with SFW and Microsoft SQL Server, see the *Storage Foundation and High Availability Solutions Quick Recovery 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 50.

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 SQL Server, 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 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 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.

These FlashSnap commands are implemented through the SFW GUI. CLI equivalents of these GUI commands exist, and are available as parameters to the `vxassist` or `vxsnap` 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 CLI `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.

The `vxsnap` command 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 SQL Server.

For more information on how VSS and SFW work together, see the *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.

See "Creating initial snapshots" on page 44.

See "Refreshing a snapshot" on page 44.

See "Recovering a database" on page 45.

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 SQL Server, 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.

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 `bpend_notify.bat` file in NetBackup or from a batch file in a pre/post command to run at the completion of a 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 by 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.

Storage Foundation 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 *Storage Foundation and High Availability Solutions Quick Recovery 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 Server, 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 *Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft SQL* for detailed procedures on Quick Recovery in a Microsoft SQL Server 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:

- See "Off-host backups" on page 46.
- See "Reporting and analysis" on page 47.
- See "Application testing and training" on page 48.

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

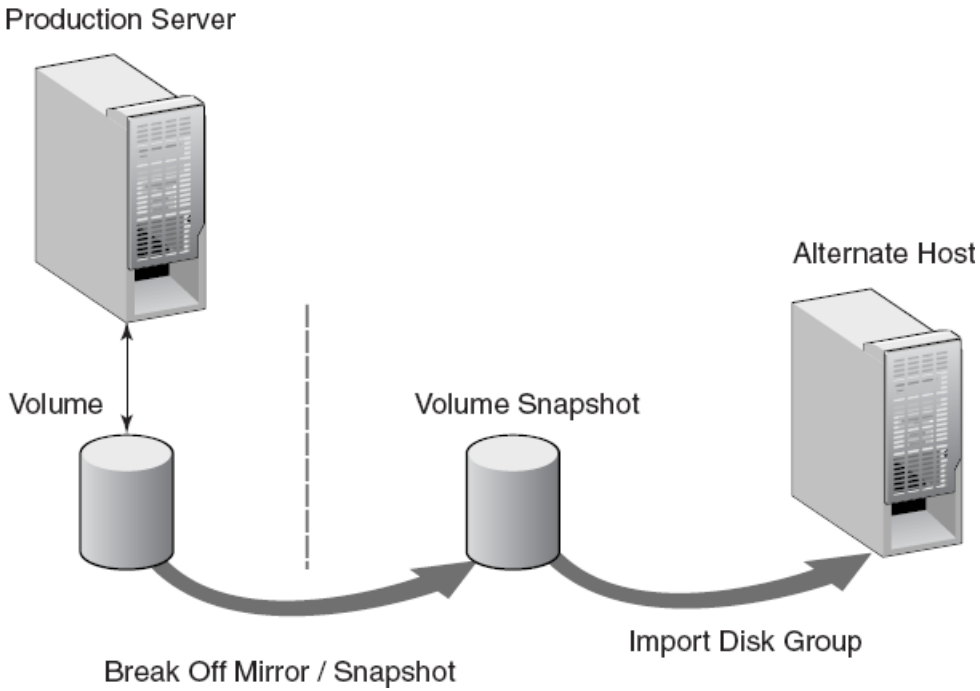

FlashSnap simplifies the process of making snapshot volumes available for off-host processing with the Disk Group Split and Join feature. Using this feature,

administrators can split one or more volume snapshots into another disk group, then “deport” the disk group. The alternate host, running Storage Foundation, 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 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

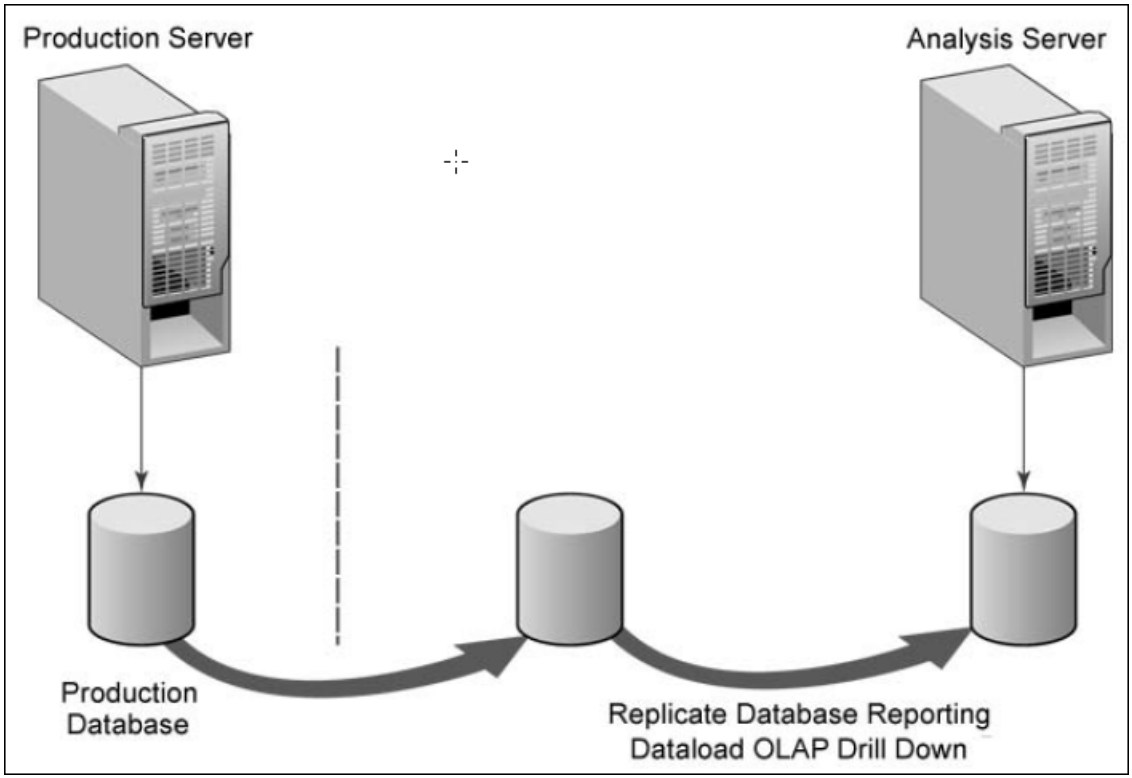

Because taking the snapshot itself has a very brief, limited impact on the production system, you can generate fresh data for analysis on a regular basis. You can even create a replica of the production database on a secondary system to be accessible for “drill-down” analysis from OLAP applications. 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.

Chapter 5

Quick Recovery example

This chapter includes the following topics:

- Example of quick recovery of an Oracle database
- Tips and references about FlashSnap

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).

Prerequisites for performing quick recovery of an Oracle database

This example assumes experience with Oracle database backup and recovery and SFW FlashSnap procedures.

The setup must be as follows:

- The Oracle database must be running in the ARCHIVELOG mode on a single server.
- The volume that contains the datafile for the Oracle database must not be the system or boot volume, but an SFW volume.

Tasks for performing quick recovery of an Oracle database

The tasks necessary for performing quick recovery in this example are:

1. Create a snapshot of the Oracle datafile volume, and resume normal processing with the Oracle datafile.

   See "Creating a split-mirror snapshot of the database" on page 51.

2. Simulate Oracle datafile corruption, and recover the Oracle datafile.

   See "Recovering the database using the split-mirror snapshot and database logs" on page 52.

Creating a split-mirror snapshot of the database

The following figure 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

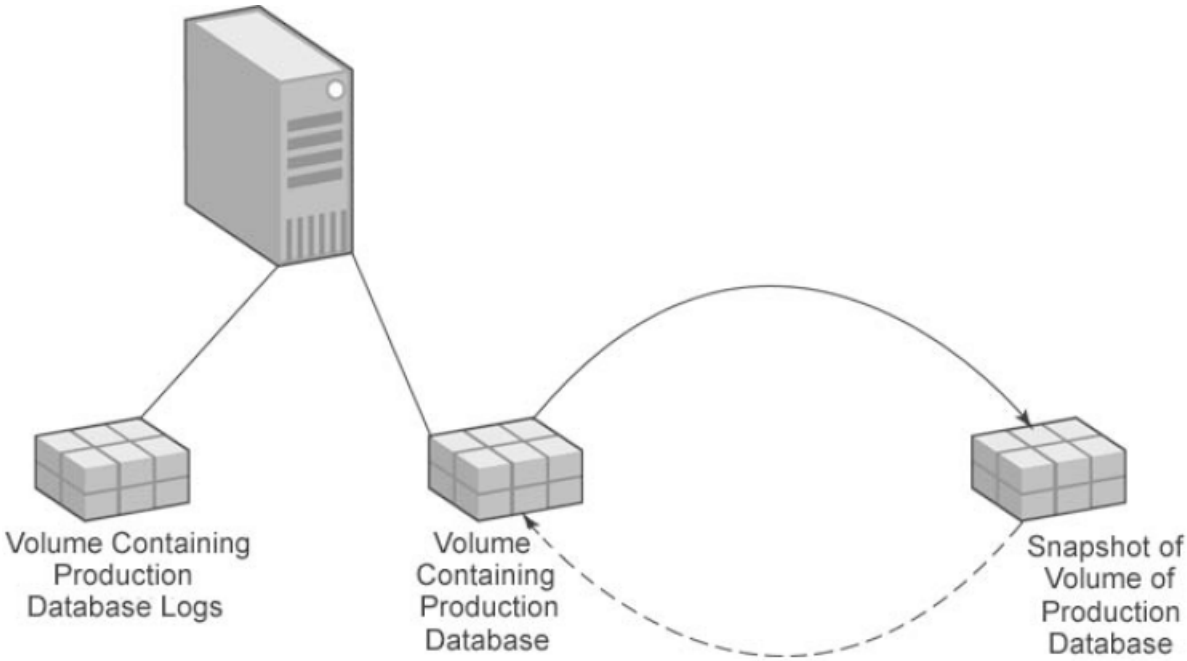

The following procedure describes how to create a snapshot of the database and save it for later use. The snapshot process is most beneficial when done on a periodic basis.

To create a snapshot of the Oracle datafile volume

1. Open the Oracle database and verify that the tablespace you want to work with is running normally.
2. In the VEAGUI, prepare the volume that contains the datafile of the tablespace.

   Optionally, use the following command:

   ```
   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 the VEA GUI, click the **Snap Shot** button to create a snapshot of the Oracle datafile volume.

   Optionally, use the following command:

   ```
   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.

To resume normal processing with the Oracle datafile

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

Recovering the database using the split-mirror snapshot and database logs

The following figure 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

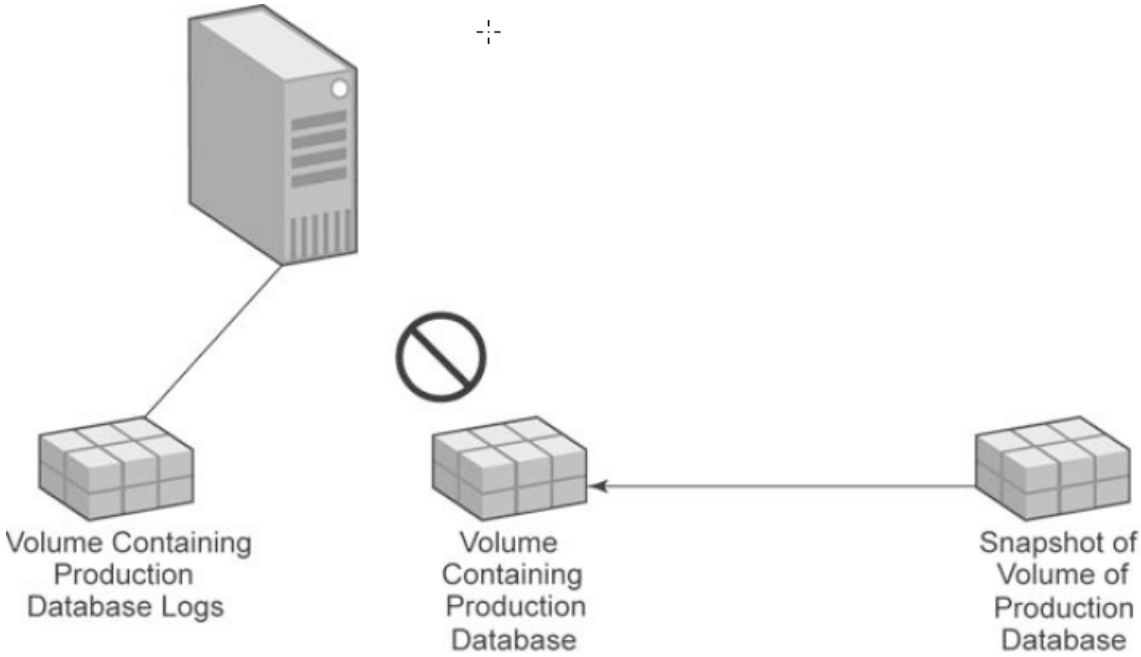

Note: This example uses a single snapshot of the datafile of the tablespace of an Oracle database. It could also be done by using multiple, simultaneous snapshots that include both the data file and the log.

To simulate Oracle datafile corruption

1. In Oracle Enterprise Manager, take the tablespace offline.
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. In the VEA GUI, click the **Snap Back** button to reattach the snapshot volume, and 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.

Tips and references about FlashSnap

Consider the following when using FlashSnap:

- 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, and so on.
- For more information on FlashSnap, see the *Storage Foundation Administrator’s Guide*.
- For more information on Quick Recovery for an application (such as SQL Server), refer to the application-specific Quick Recovery Solutions guide.

Section 3

High Availability

- Chapter 6. High availability: Overview
- Chapter 7. Deploying InfoScale Enterprise for high availability: New installation
- Chapter 8. Adding DMP to a clustering configuration

Chapter 6

High availability: Overview

This chapter includes the following topics:

- About high availability
- About clusters
- How VCS monitors storage components

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 Cluster Server (VCS) with Storage Foundation.

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.

How VCS monitors storage components

VCS provides specific agents that monitor storage components and ensure that the shared disks, disk groups, LUNs, volumes, and mounts are accessible on the system where the application is running. Separate agents are available for shared and non-shared storage and for third-party storage arrays such as NetApp filers. Your storage configuration determines which agent should be used in the high availability configuration.

For details on the various VCS storage agents, refer to the *Cluster Server Bundled Agents Reference Guide*.

Shared storage—if you use NetApp filers

The VCS hardware replication agents for NetApp provide failover support and recovery in environments that employ NetApp filers for storage and NetApp SnapMirror for replication. The agents enable configuring NetApp filers over an iSCSI or Fibre Channel (FC) connection in a VCS cluster environment.

The VCS agents for NetApp are as follows:

- NetAppFiler
- NetAppSnapDrive
- NetAppSnapMirror

These agents monitor and manage the state of replicated filer devices and ensure that only one system has safe and exclusive access to the configured devices at a time. The agents can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments that are set up using the VCS Global Cluster Option (GCO).

In a typical configuration, the agents are installed on each system in the cluster. The systems are connected to the NetApp filers through a dedicated (private) storage network. VCS cluster systems are physically attached to the NetApp filer via an ethernet cable supporting iSCSI or FC as the transport protocol.

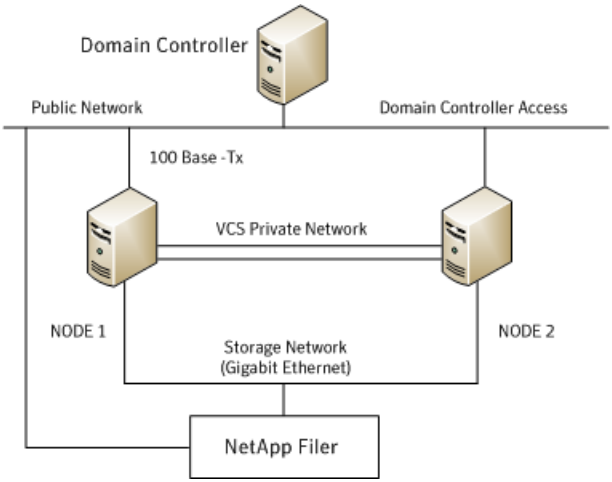

VCS also provides agents for other third-party hardware arrays. For details on the supported arrays, refer to the product Software Compatibility List (SCL).

Shared storage—if you use SFW to manage cluster dynamic disk groups

The VCS MountV and VMDg agents are used to monitor shared storage that is managed using Storage Foundation (SFW). SFW manages storage by creating disk groups from physical disks. These disk groups are further divided into volumes that are mounted on the cluster systems.

The MountV agent monitors volumes residing on disk groups. The VMDg agent monitors cluster dynamic disk groups and is designed to work using SCSI reservations. Together the MountV and VMDg agents ensure that the shared cluster dynamic disk groups and volumes are available.

Shared storage—if you use Windows LDM to manage shared disks

The VCS Mount and DiskReservation (DiskRes) agents are used to monitor shared disks that are managed using Windows Logical Disk Management (LDM).

The Mount agent monitors basic disks and mount points and ensures that each system is able to access the volume or mount path in the same way. The DiskRes agent monitors shared disks and uses persistent reservation to ensure that only one system has exclusive access to the disks. During failovers, these agents ensure that the disks and volumes are deported and imported on the node where the application is running.

Non-shared storage—if you use SFW to manage dynamic disk groups

VCS introduces the Volume Manager Non-Shared Diskgroup (VMNSDg) agent to support local non-shared storage configurations that are managed using SFW. The VMNSDg agent works without SCSI reservations and is designed for locally attached storage devices that do not support SCSI.

The VMNSDg agent monitors and manages the import and deport of dynamic disk groups created on local storage. The only difference between the VMDg agent and the VMNSDg agent is that the VMDg agent is designed for shared cluster dynamic disk groups and uses SCSI reservations, whereas the VMNSDg agent supports only non-shared local dynamic disk groups and works without SCSI reservations.

The VMNSDg agent can be used to set up single node Replicated Data Clusters (RDC) or Disaster Recovery (DR) configurations with replication set up between the sites.

During a failover, the VCS MountV and VMNSDg agents deport the locally attached storage from the affected node and then import the locally attached storage of the target node. Replication ensures that the data is consistent and the application is up and running successfully.

Note: The VMNSDg agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use Windows LDM to manage local disks

VCS introduces the NativeDisks agent to support local non-shared storage configurations managed using Windows LDM. The NativeDisks agent works without SCSI reservations and is designed for local storage that does not support SCSI.

Together with the Mount agent, the NativeDisks agent monitors and manages the import and deport of basic local disks on the system. The only difference between the DiskRes agent and the NativeDisks agent is that the DiskRes agent is designed for shared disks and uses SCSI reservations, whereas the NativeDisks agent supports only non-shared local disks and works without SCSI reservations.

Note: The NativeDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use VMware storage

VCS introduces the VMwareDisks agent to support storage configurations in a VMware virtual environment. The agent is platform independent and supports VMware Virtual Machine Disk (VMDK), Raw Device Mapping (RDM) disk files (virtual), and storage that is configured using Network File System (NFS). The VMwareDisks agent works without SCSI reservations and supports locally attached non-shared storage.

VMware features such as snapshots, vMotion, and DRS do not work when SCSI disks are shared between virtual machines. The VMwareDisks agent is designed to address this limitation. With this agent, the disks can now be attached to a single virtual machine at a time in the VCS cluster. On failover, along with the service group, the VMwareDisks agent moves the disks to the target virtual machine.

The VMwareDisks agent communicates with the host ESXi server to configure storage. This agent manages the disk attach and detach operations on a virtual machine in the VCS cluster. The agent is VMware HA aware. During failovers, the agent detaches the disk from one system and then attaches it to the system where the application is actively running. The VMwareDisks agent presents the virtual disks to the operating system. On Windows, the agent relies on the VMNSDg agent (in case of SFW-managed local storage) and the NativeDisks agent (in case of LDM-managed local storage) for initializing and managing the virtual disks. On Linux, the agent relies on the LVM and VxVM agents.

Note: The VMwareDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Chapter 7

Deploying InfoScale Enterprise for high availability: New installation

This chapter includes the following topics:

- About the high availability solution
- Tasks for a new high availability (HA) installation—additional applications
- Reviewing the InfoScale installation requirements
- Notes and recommendations for cluster and application configuration
- Reviewing the configuration
- Configuring the storage hardware and network
- About installing the Veritas InfoScale products
- Configuring disk groups and volumes
- Configuring the cluster using the Cluster Configuration Wizard
- About modifying the cluster configuration
- About installing and configuring the application or server role
- Configuring the service group
- Configuring the service group in a non-shared storage environment

- Verifying the cluster configuration
- Possible tasks after completing the configuration
- Modifying the application service groups

About the high availability solution

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

Veritas recommends using the Solutions Configuration Center as a guide for setting up high availability with InfoScale Enterprise.

See "About the Solutions Configuration Center" on page 23.

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.

The following table outlines the high-level objectives for implementing the configuration and the tasks for each objective.

Table 7-1 Task list: New High Availability configuration

Objectives	Tasks
See "Reviewing the InfoScale installation requirements" on page 64.	■ Verify the hardware and software prerequisites. ■ Review the requirements.
See "Reviewing the configuration" on page 67.	Review the configuration
See "Configuring the storage hardware and network" on page 335.	Configure the storage hardware and network.
See "About installing the Veritas InfoScale products" on page 503.	Install the product.
See "Configuring disk groups and volumes" on page 70.	■ Plan your storage layout. ■ Create disk groups. ■ Create volumes. ■ Manage disk groups and volumes.
See "Configuring the cluster using the Cluster Configuration Wizard" on page 345.	Use the VCS Cluster Configuration Wizard (VCW) to set up the cluster
See "About installing and configuring the application or server role" on page 108.	■ As necessary, install the application program files on the local system drive of the first node. ■ Install files relating to the data and logs on the target storage. ■ Deport the disk groups on the first node and import them on the second node. Install the application on the second node.
See "Configuring the service group" on page 110.	■ Use the applicable wizard to create and configure the VCS service group or groups. ■ Create the application service group manually using templates from the Cluster Manager (Java Console) (if using a non-shared storage configuration). ■ Bring the service group online.
See "Enabling fast failover for disk groups (optional)" on page 162.	■ Use the Java Console to enable the FastFailover attribute for VMDg resources.
See "Verifying the cluster configuration" on page 269.	■ Switch the service group to the second node. ■ Shut down an active cluster node.

Table 7-1 Task list: New High Availability configuration *(continued)*

Objectives	Tasks
See "Possible tasks after completing the configuration" on page 166.	■ Modify the cluster configuration. ■ Modify the application or server role service group.

Reviewing the InfoScale installation requirements

Before installing InfoScale Enterprise, review the product installation requirements for your systems. Minimum requirements and Veritas-recommended requirements may vary. For details, see the *Veritas InfoScale Installation and Upgrade Guide.*

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
  https://sort.veritas.com/documents

  Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

  The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

  Refer to the following Microsoft knowledge database article for more details:

  http://technet.microsoft.com/en-us/library/dd184075.aspx

- 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 *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. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

  Static IP addresses are required for the following purposes:

  - One static IP address 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 Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
  - For Volume Replicator 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 SFW HA 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 *Cluster Server Bundled Agents Reference Guide*.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the

network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.

- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- 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.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.
- For a Replicated Data Cluster, install only in a single domain.
- 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)
- This is applicable for a Replicated Data Cluster configuration.
  This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, 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.
- To configure a RDC cluster, you need to 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.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.
VCS agents, wizards, and other components	VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above). The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

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 “About Dynamic Multi-Pathing” on page 177.

Figure 7-1 SFW HA Active-Passive configuration with two servers

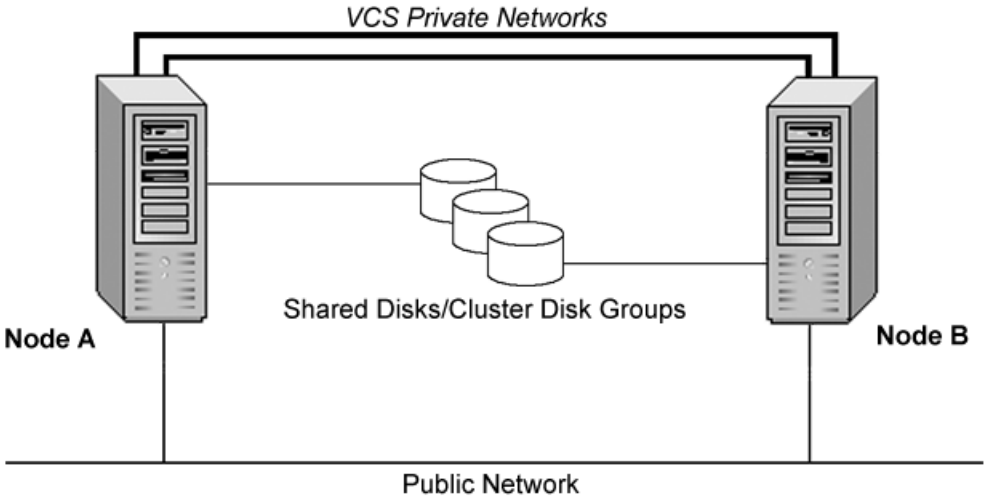

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

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, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
   - Veritas 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.

To verify the DNS settings and binding order for all systems

1. Open the Control Panel by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.
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, click **Adapter settings**.
4. Ensure the public network adapter is the first bound adapter by following these steps sequentially:
   - In the Network Connections window, click **Advanced > 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)** or **Internet Protocol Version 6 (TCP/IPv6)** check box, depending on which protocol your network is using.
   - 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 that the **Register this connection's address in DNS** check box is selected.

12 Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.

13 Click **OK**.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Configuring disk groups and volumes

Use Storage Foundation 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 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 *Storage Foundation Administrator's Guide* for more information.

Configuring disk groups and volumes involves the following tasks:

- See "Planning disk groups and volumes" on page 71.
- See "Creating a dynamic disk group" on page 508.
- See "Creating dynamic volumes" on page 475.

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:

- See "Planning your File Share storage" on page 71.
- See "Planning your Print Share storage" on page 71.
- See "Planning your IIS storage" on page 71.
- See "Planning your storage for additional applications" on page 72.
- See "Considerations for a fast failover configuration" on page 73.

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 (VMDg) or non-shared storage (VMNSDg).
- When configuring a new set up, first create the disk groups and volumes on the shared storage (VMDg) or non-shared storage (VMNSDg) and then create the directory structure for the file shares.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage (VMDg) or non-shared storage (VMNSDg) using the practices recommended by Microsoft.

Planning your Print Share storage

Make sure you create separate volumes on shared storage for the print spooler and for the RegRep volume that contains the list of registry keys that must be replicated between the cluster systems.

Note: Veritas recommends that you place the RegRep and spooler data in separate directories on separate volumes on the shared storage.

Note: The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems.

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 (VMDg) or non-shared storage (VMNSDg).
- For a new IIS installation, make sure that the directory for the web sites is created on volumes on the shared storage (VMDg) or non-shared storage (VMNSDg).
- For existing web sites, stop the sites and then move the website content to volumes on the shared storage (VMDg) or non-shared storage (VMNSDg). You must also reconfigure the home directory location for the web site in IIS and then restart the web site again.

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 Veritas 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 system drive of the server. Data files and other related files, such as logs, are placed on the shared storage (VMDg) or non-shared storage (VMNSDg). 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 Volume Replicator, 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
- The type of replication that you plan to use
  If you plan to implement a disaster recovery configuration with Volume Replicator, 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.

Considerations for a fast failover configuration

For VCS service groups that contain many disk groups, you can greatly reduce failover time by implementing fast failover.

Fast failover speeds up the failover of storage resources in several ways:

- Fast failover provides a "read-only deported" mode for disk groups on inactive nodes. This mode speeds up the process of importing a disk group.
- Fast failover maintains the current disk group configuration in memory on the inactive nodes. Any changes are automatically synchronized so that all nodes maintain an identical disk group configuration.

For more details about fast failover, refer to the *Storage Foundation Administrator's Guide*.

Take the following storage-related requirements into account if you are planning to implement fast failover:

- Fast failover is currently not supported for the following:
  - RAID-5 volumes
  - SCSI-2
  - Active/Passive (A/P) arrays for DMP
- In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS and ReFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode.
- The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click Properties. Disk group version upgrade is required after upgrading SFW HA on

the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

Creating a dynamic disk group

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.

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

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

To create a dynamic disk group

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:
   - Enter the disk group name (for example, **DG1**).
   - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.

- 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.
- Click **Next**.

Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

---

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the dynamic 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.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

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

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. For Site Preference, leave the setting as **Siteless** (the default).
- Automatic disk selection is the default setting. To manually select the disks, click **Manually select disks** 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.
- Click **Next**.

7 Specify the volume attributes.

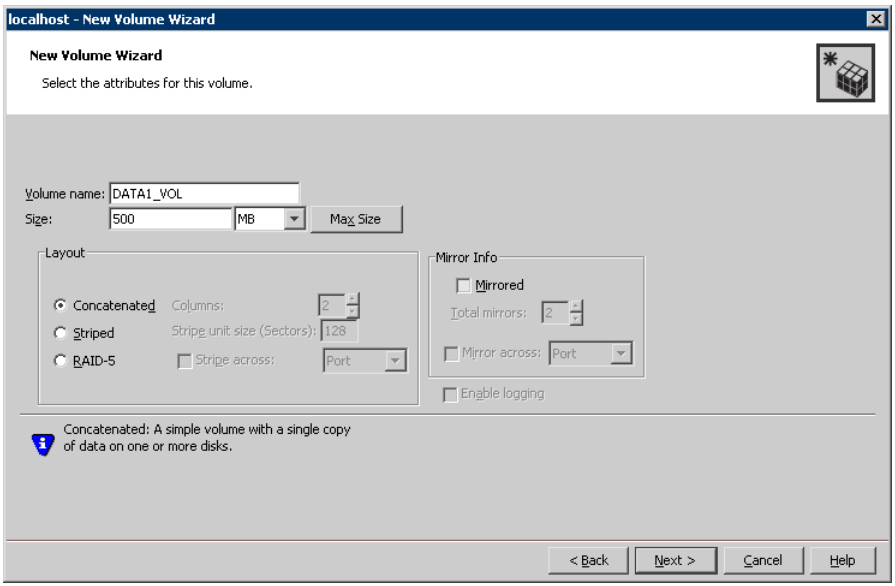

- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns** 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.
- Click **Next**.

8 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 disk.
- Click **Next**.

9 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**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

About managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.

- To move a cluster dynamic 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.

Managing disk groups and volumes involves the following:

- See "Importing a disk group and mounting a volume" on page 337.
- See "Unmounting a volume and deporting a disk group" on page 338.

Note: (Disaster recovery configurations only) 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 *Storage Foundation Administrator's Guide* for more information.

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 using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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 VCW to remove the node from the cluster, rename the system, and then run VCW again to add that system to the cluster.

Note the following prerequisites before you proceed:

- The required network adapters, and SCSI controllers are installed and connected to each system.

  To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.

- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

  Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list.

  When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.

Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.

- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

1. Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
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.

   To discover information about all systems and users in the domain, do the following:

   - Clear **Specify systems and users manually**.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - Select **Specify systems and users manually**.
     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. 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.

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**.

Select the 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.

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.
- Product is either not installed or there is a version mismatch.

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

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

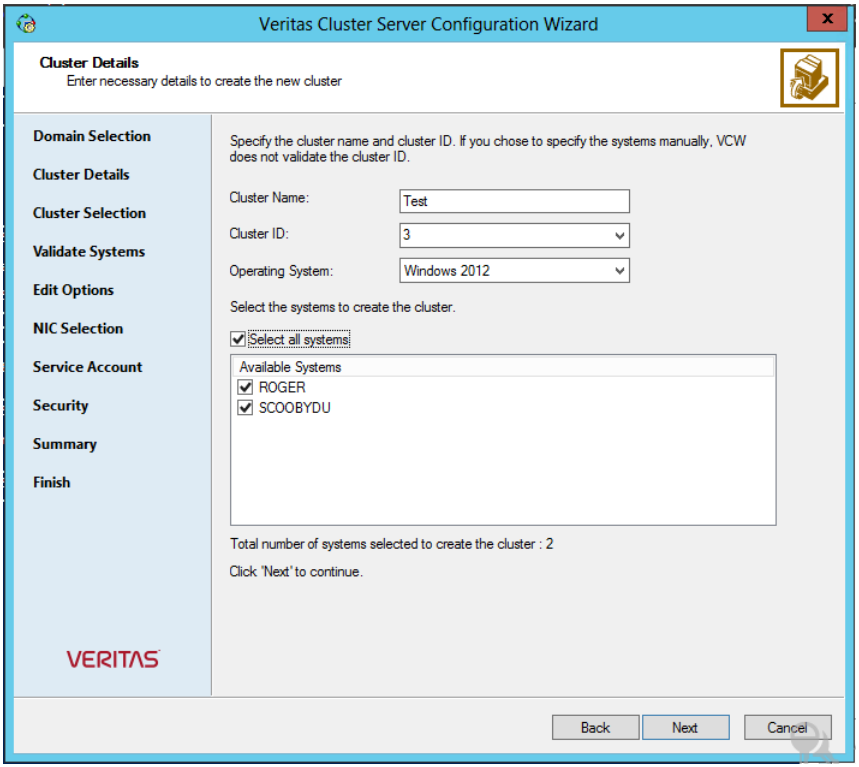

Specify the cluster details as follows:

Cluster Name	Type a name for the new cluster. Veritas 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 65535. **Note:** 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. All the systems in the cluster must have the same operating system and architecture. For example, you cannot configure a Windows Server 2008 R2 system and a Windows Server 2012 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 step 9, proceed to the next step. Otherwise, proceed to step 12.

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 using IPv4 or IPv6 network.

Do one of the following:

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

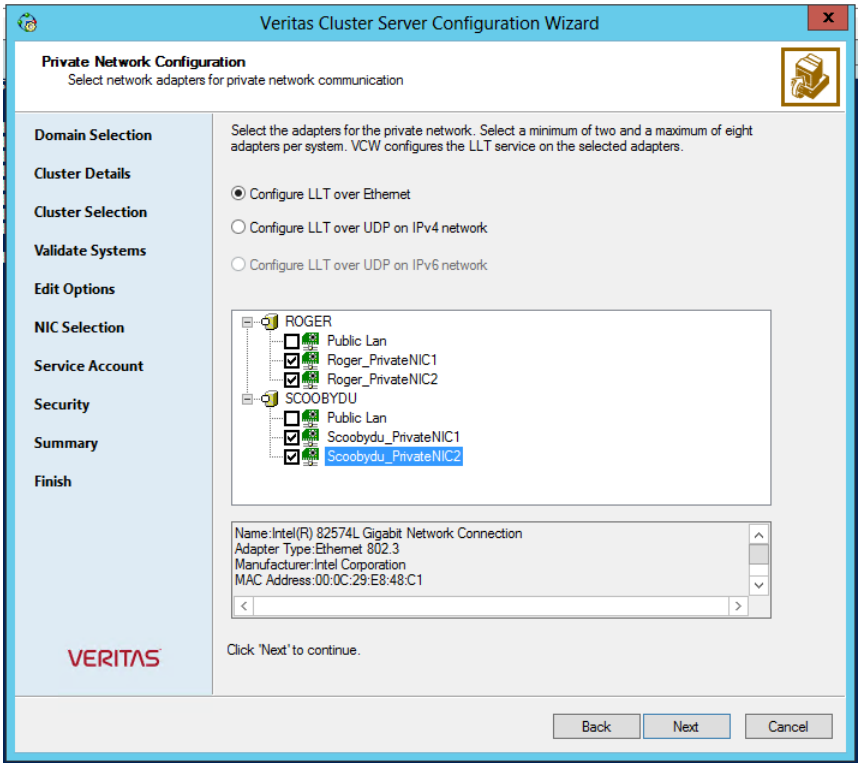

- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.
  Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends 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. Veritas 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 on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.
  The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the VCS private network.
- 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. In case of IPv4, each IP address can be in a different subnet.
  The IP address is used for the VCS private communication over the specified UDP port.
- 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, 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 are used for the private network communication.

12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service.

The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details 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, 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 communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, 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 Cluster Manager (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.

  Veritas recommends that you specify a new user name and password.

- To use the single sign-on feature, click **Use Single Sign-on**.

  In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.

  The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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 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.

16 On the Cluster Service Components panel, select the components to be configured in the ClusterService group and then click **Next**.

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 354.
- 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.

  Note: Configure GCO using this wizard only if you are configuring for disaster recovery but not using the DR wizard.

  Alternatively, consider configuring GCO using the DR wizard.

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 then 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 then click **Next**.

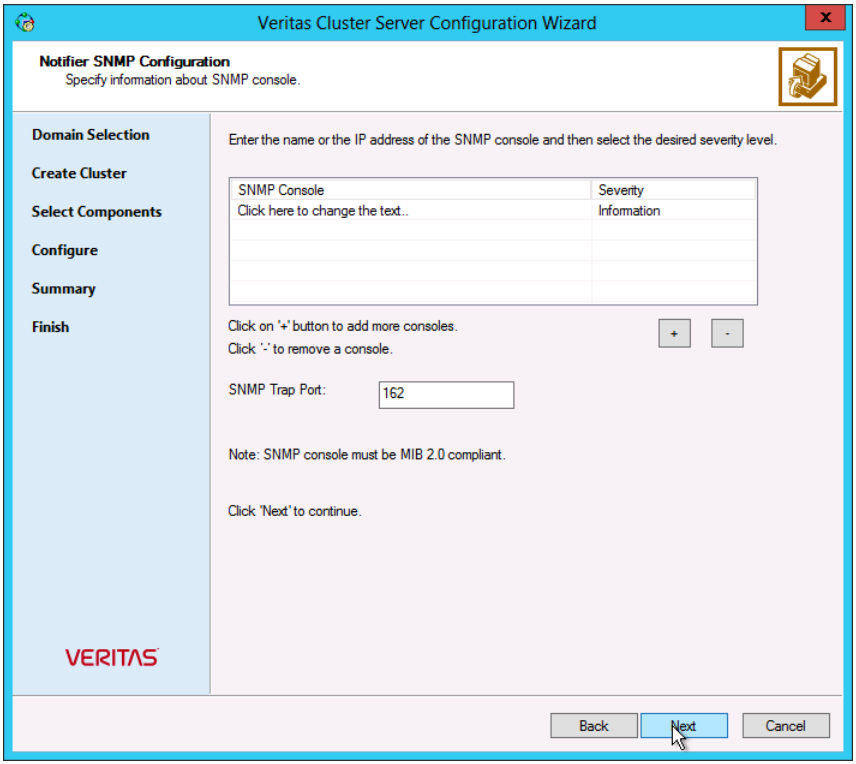

Do the following:

- 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 the **+** icon to add a field; click the **-** icon 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 then click **Next**.

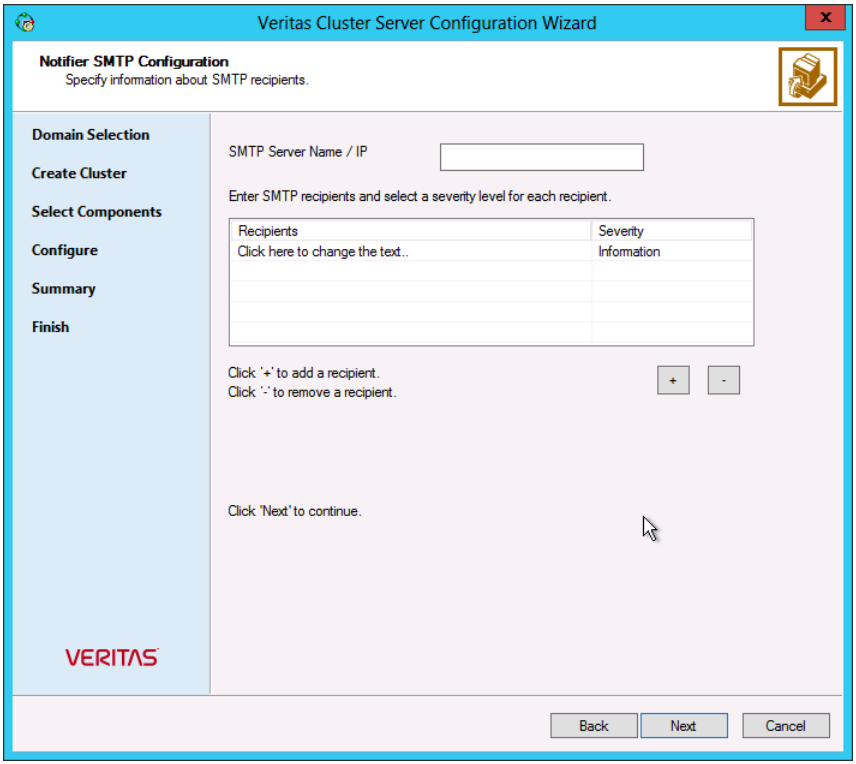

Do the following:

- 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 the **+** icon to add fields; click the **-** icon to remove a field.

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

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.

6 Click **Finish** to exit the wizard.

About modifying the cluster configuration

You can modify a cluster configuration using the VCS Cluster Configuration Wizard (VCW).

When used to modify a cluster configuration, the wizard lets you perform the following tasks:

- Adds nodes to a cluster
  See "Adding nodes to a cluster" on page 271.
- Remove nodes from a cluster
  See "Removing nodes from a cluster" on page 96.
- Reconfigure the private network and LLT
  See "Reconfiguring a cluster" on page 98.
- Configure the ClusterService service group in a cluster
  See "Configuring the ClusterService group" on page 102.
- Delete a cluster

  Note: When used to delete a cluster configuration, the wizard removes the cluster components from the nodes; it does not uninstall the VCS components.

  See "Deleting a cluster configuration" on page 106.

You can perform the following modifications on a cluster manually:

- Configure single sign-on for the cluster
- Convert a secure cluster to a non-secure cluster

Adding nodes to a cluster

If you are setting up a Replicated Data Cluster, use the VCS Cluster Configuration Wizard (VCW) to add the systems in the secondary zone (zone1) to the existing cluster.

You use the VCS Cluster Configuration Wizard (VCW) to add one or more nodes to an existing cluster.

Prerequisites for adding a node to an existing cluster are as follows:

- Verify that the logged-on user has VCS cluster administrator privileges.
- The logged-on user must be a local administrator on the system where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select Services on the Administrative Tools menu and verify that the Veritas Command Server shows that it is started.
- Verify that the high availability daemon (HAD) is running on the node on which you run the wizard. Open the Services window, and verify that the **Veritas High availability engine** service is running.

To add a node to a VCS cluster

1. Start the VCS Cluster Configuration wizard.

   Click **Start > All Programs > Veritas > 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.

   To discover information about all the systems and users in the domain, do the following:

   - Clear the **Specify systems and users manually** check box.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - 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. Otherwise proceed to the next step.

5 On the System Selection panel, complete the following and click **Next**:

- Type the name of an existing 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.

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.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's 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.

Click **Next** to proceed.

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 (a cluster that is not a secure cluster).

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 use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over Ethernet, do the following:
  - Select the check boxes next to the two NICs to be assigned to the private network.

    Veritas 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 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. Veritas 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, do the following:
  - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Veritas recommends reserving at least two NICs exclusively for the VCS private network. 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.
  - 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. In case of IPv4, 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 are 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 credentials 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**.

If you are setting up a Replicated Data Cluster, return to the task list:

See "Creating a parallel environment in the secondary zone" on page 270.

Removing nodes from a cluster

This topic describes how to remove nodes from a multiple node VCS cluster. To remove a node from a single node cluster, you must delete the cluster.

See "Deleting a cluster configuration" on page 106.

To remove nodes from a cluster

1 Verify that no service groups are online on the node to be removed.

2 Remove the node from the SystemList of all service groups.

3 Start the VCS Configuration wizard.

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

4 Read the information on the Welcome panel and click **Next**.

5 In the Configuration Options panel, click **Cluster Operations** and click **Next**.

6 In the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the domain discovery options.

   To discover information about all the systems and users in the domain:

   - Uncheck the **Specify systems and users manually** check box.
   - Click **Next**.
   - Proceed to step 10.

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 checked **Retrieve system list from domain**, proceed to step 8. Otherwise proceed to the next step.

7 In the System Selection panel, type the name of the system and click **Add**.

Proceed to step 10.

8 In the System Selection panel, specify the systems for the cluster from which you will be removing the nodes.

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.

9 The System Report panel displays the validation status, whether Accepted or Rejected, of all the systems you specified earlier.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's 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.

Click **Next** to proceed.

10 In the Cluster Configuration Options panel, click **Edit Existing Cluster** and then click **Next**.

11 In the Cluster Selection panel, select the cluster to be edited and click **Next**.

If you chose to specify the systems manually in step 6, only the clusters configured with the specified systems are displayed.

12 In the Edit Cluster Options panel, click **Remove Nodes** and then click **Next**.

In the Cluster User Information panel, enter 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 remove a node from a non-secure cluster.

13 In the Cluster Details panel, select the check boxes next to the nodes to be removed and click **Next**.

See "Reconfiguring a cluster" on page 98.

14 If you want to remove the VCS Helper service user account from the administrative group of the nodes being removed from the cluster, click **Yes** from the informational dialog box. Otherwise, click **No**.

15 The wizard validates the selected nodes. 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.

An informational dialog box appears if you are removing all but one nodes of a multiple node cluster. In the dialog box, specify whether you want to retain or remove the private link heartbeat.

16 Review the summary information and click **Remove**.

The wizard starts running commands to remove the node from the cluster.

17 After the commands have been successfully run, click **Finish**.

Reconfiguring a cluster

You may need to reconfigure your cluster after changing an adapter on a cluster node, to update the LLT information, or to configure Veritas Security Services.

To reconfigure a cluster

1 Start the VCS Configuration wizard.

Click **Start > All Programs > Veritas > Veritas Cluster Server> Configuration Tools > Cluster Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu.

2 Read the information on the Welcome panel and click **Next**.

3 In 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 click **Next**.

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
- Proceed to step 8.

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 checked **Retrieve system list from domain**, proceed to step 6. Otherwise proceed to the next step.

5 In the System Selection panel, type the name of the system and click **Add**.

Proceed to step 8.

6 In the System Selection panel, specify the systems for the cluster to be reconfigured.

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.

7 The System Report panel displays the validation status, whether Accepted or Rejected, of all the systems you specified earlier.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's 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.

Click **Next** to proceed.

8 In the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.

9 In the Cluster Selection panel, select the cluster to be reconfigured 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 In the Edit Cluster Options panel, click **Reconfigure** and click **Next**.

In the Cluster User Information dialog box, enter 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 reconfigure a non-secure cluster.

11 In the second Edit Cluster Options dialog box, select any of the following options and click **Next**:

- **Change private network heartbeat links**
  Select this option to change the private network heartbeat links. If the selected cluster is a single node cluster, the option is to remove the private heartbeat links.
  If the cluster has more than one node, the options are to add or remove private heartbeat links.
  See step 12.
- **Change HAD Helper User account**
  Selection this options to change the user account for the VCS Helper service.
  See step 13.
- **Configure VCS Authentication Service**
  Select this option to configure the VCS authentication service for single sign-on. Single sign-on configures a secure cluster.

12 If the option to change the private network heartbeat links was selected, do one of the following:

- To configure the VCS private network over Ethernet, do the following:
  - Select the check boxes next to the two NICs to be assigned to the private network.
    Veritas 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 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. Veritas 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, do the following:
  - Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on how LLT is configured on the existing nodes in the cluster.
  - Select the check boxes next to the NICs to be assigned to the private network. You can assign maximum eight network links.
    Veritas recommends reserving two NICs exclusively for the VCS private network.
  - 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.
  - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. The default ports numbers are 50000 to 50007. You can use ports in the range 49152 to 65535. 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. In case of IPv4, 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 are used for the private network communication.

13 If the option to change the VCS HAD Helper User account was selected, in the VCS Helper Service User Account dialog box, specify the name of a domain user in whose context the VCS Helper service will run.

The VCS High Availability Daemon, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

- Select one of the following:
  - **Existing user**
    Choose an existing user account context for the VCS Helper service.
  - **New user**
    Create a new user account context for the VCS Helper service.
- Enter a valid user name for the selected account and click **Next**.

Do not append the domain name to the user name; do not enter user names as `DOMAIN\user` or `user@DOMAIN`.

- Enter a password for the selected account and click **OK**.

14 Review the summary information and click **Reconfigure**.

15 The wizard starts running commands to apply the changes. After all services have been successfully configured, click **Finish**.

Configuring the ClusterService group

Use the VCS Configuration wizard to configure the following ClusterService service group components, if you did not configure them during the initial cluster configuration:

- Notification
- GCO Option for inter-cluster communication for global clusters

Note that the wizard allows you to configure each component only once.

To configure the ClusterService group

1 Start the VCS Configuration wizard.

Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu.

2 Read the information on the Welcome panel and click **Next**.

3 In 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 click **Next**.

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 7.

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 checked the **Retrieve system list from domain** check box, proceed to step 6. Otherwise proceed to the next step.

5 In the System Selection panel, type the name of the system and click **Add**.

Proceed to step 7.

6 In the System Selection panel, specify the systems for the cluster where you will be configuring the ClusterService group.

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 will discover all the nodes for that cluster.

7 In the Cluster Configuration Options panel, click **Edit Existing Cluster** and then click **Next**.

8 In the Cluster Selection panel, select the cluster to be edited and click **Next**.

If you chose to specify the systems manually in 4, only the clusters configured with the specified systems are displayed.

9 In the Edit Cluster Options panel, click **Configure ClusterService Options** and then click **Next**.

In the Cluster User Information dialog box, enter 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 configure a ClusterService group in a non-secure cluster.

10 In the Cluster Service Components panel, select from the following components to be configured in the ClusterService service group and then click **Next**.

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 104.
- 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.
  See "Configuring the wide-area connector process for global clusters" on page 105.

Configuring notification

This topic describes how to configure the notifier resource.

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**.

Configure the SNMP console as follows:

- 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 server, specify information about SMTP recipients and click **Next**.

Configure the SMTP server as follows:

- 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**.

Specify the network information on the Notifier Network Card Selection panel as follows:

- 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 are done with the configuration, click **Finish** to exit the wizard.

Configuring the wide-area connector process for global clusters

This topic describes how to configure wide-area connector resource 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 group configured, you can use the IP address configured in the service group or configure a new IP address.

Do the following:

- To specify an existing IP address, select **Use existing IP resource** and then select the IP address from the drop-down list.
- To use a new IP address, do the following:

- In case of IPv4, select **IPV4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- In case of IPv6, select **IPV6** and select the IPv6 prefix from the drop-down list.
  The wizard uses the prefix and automatically generates a unique IPv6 address that is valid on the network. The IPv6 option is disabled if the network does not support IPv6.

- 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.

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.

The wizard does not set up a global cluster environment; it configures a resource for the wide-area connector, which is required for inter-cluster communication.

For instructions on setting up a global cluster environment:

See "Configuring the global cluster option for wide-area failover" on page 405.

Deleting a cluster configuration

This topic describes how to delete a cluster configuration.

To delete a cluster configuration

1 Start the VCS Configuration wizard.

Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu.

2 Read the information on the Welcome panel and click **Next**.

3 In 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 click **Next**.

To discover information about all the systems and users in the domain

- Uncheck the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 7.

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 checked the **Retrieve system list from domain** check box, proceed to step 6. Otherwise proceed to the next step.

5 In the System Selection panel, type the name of the system and click **Add**.

Proceed to step 7.

6 In the System Selection panel, specify the nodes of the cluster to be deleted.

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, VCW discovers all nodes for that cluster.

7 In the Cluster Configuration Options panel, click **Delete Cluster** and then click **Next**.

8 In the Cluster Selection panel, select the cluster whose configuration is to be deleted and click **Next**.

If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

9 If you want to remove the VCS Helper service user account from the administrative group of the all the nodes in the cluster, click **Yes** from the informational dialog box. Otherwise, click **No**.

10 In the Cluster User Information panel, enter 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 delete a non-secure cluster.

11 Review the summary information and click **Unconfigure**.

12 The wizard starts running commands to remove the configuration from the cluster. After all commands have been successfully run, click **Finish**.

VCW removes the cluster configuration; VCW does not unconfigure the VCS Authentication Service or uninstall the product from the systems.

About 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:

- See "About configuring a File Share server role" on page 108.
- See "About configuring a Print Share server role" on page 108.
- See "About installing and configuring the IIS application" on page 109.
- See "About installing additional applications" on page 109.

About 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 (VMDg) or non-shared storage (VMNSDg).
- When configuring a new set up, first create the disk groups and volumes on the shared storage (VMDg) or non-shared storage (VMNSDg) and then create the directory structure for the file shares.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage (VMDg) or non-shared storage (VMNSDg) using the practices recommended by Microsoft.
- The FileShare agent is installed automatically with InfoScale Enterprise.

About 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.

Note: The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems.

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.

About installing and configuring the IIS application

Points to note when installing IIS:

- Verify that IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on the shared storage (VMDg) or non-shared storage (VMNSDg).
- Import the 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 (VMDg) or non-shared storage (VMNSDg) and place the site content in that folder.
- Change the default home directory path for all the IIS sites to be monitored to a location on the shared storage (VMDg) or non-shared storage (VMNSDg). 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 (VMDg) or non-shared storage (VMNSDg). 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.

About 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 (VMDg) or non-shared storage (VMNSDg).

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:

- See "About configuring file shares" on page 110.
- See "About configuring print shares" on page 124.
- See "About configuring IIS sites" on page 132.
- See "Modifying an application service group" on page 170.
- See "About configuring the Oracle service group using the wizard " on page 156.
- See "Enabling fast failover for disk groups (optional)" on page 162.

About configuring file shares

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.

VCS provides several ways to configure file shares, including the configuration wizard, Cluster Manager (Java Console), and the command line. This section provides instructions on how to use the File Share Configuration Wizard to configure file shares.

On Windows 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 VCS commands from the command line, or remotely using the Cluster Manager (Java Console).

If you want to configure file shares manually, consider the following:

- To configure a shared directory, use the FileShare agent.
- To configure multiple directories, use the CompositeFileShare agent.
- If UAC is enabled, run the program or commands in the “Run as administrator” mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.
- Before configuring the service group, review the agent resource types and the attribute definitions described in the *Cluster Server Bundled Agents Reference Guide*.

Before you configure a file share service group

Note the following prerequisites before you configure a file share service group:

- Verify that you have local administrator privileges on the system from where you run the wizard.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.
  For a detailed list of services and ports used, refer to the product 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 disks that are accessible from the nodes that will be part of the file share 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). VEA is available with SFW HA only.
- Mount the drives or LUNs containing the shared directories on the system where you run the wizard. Unmount the drives or LUNs from other systems in the cluster.
- Verify that the Veritas Command Server service is running on all the systems in the cluster.
- If NetBIOS is disabled over TCP/IP, you must set the Lanman agent's DNSUpdateRequired attribute value to 1 (True).
  You can modify the Lanman resource attribute value after configuring the service group.

- Verify that you have the following information ready. The wizard prompts you for these details:
  - A unique virtual computer name to be assigned to the file share server
    This is the name that the clients use to access the file shares. The virtual name must not exceed 15 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
    The virtual IP address is required only if you wish to configure an IPv4 address. In case of IPv6, the wizard prompts you to select the IPv6 network and automatically generates an IPv6 address that is valid and unique on the network. The wizard uses the prefix that is advertised by the router on the IPv6 network.

    Note: Windows Server does not support accessing file shares using a virtual IP address. You can work around this restriction by using non-scoped file shares.

    See "Creating non-scoped file shares configured with VCS " on page 119.

    See "Making non-scoped file shares accessible while using virtual server name or IP address if NetBIOS and WINS are disabled" on page 121.

  - The list of directories to be shared.
    You can 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.

Configuring file shares using the wizard

The File Share Configuration Wizard enables you to create and modify file share service groups, making file shares highly available in a VCS cluster.

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.

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See "Configuring the service group in a non-shared storage environment" on page 163.

To configure file shares using the File Share Configuration Wizard

1. Start the File Share Configuration Wizard.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC). In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **File Share Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on the Welcome panel and click **Next**.

3. On the Wizard Options panel, click **Create service group** and then click **Next**.

4. On the Service Group Configuration panel, specify the following service group details:

Service Group Name	Type a name for the file share service group.
Group System List	Specify the systems on which to configure the service group. To add systems to the service group's system list, select the systems in the **Available Cluster Systems** list and click the right arrow. To remove systems from the service group's system list, select the systems 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 from the **Systems in Priority Order** and click the up and down arrow. 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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox.

   Click **Next**.

5 On the FileShare Configuration panel, specify the following configuration information for the file share resources to be created.

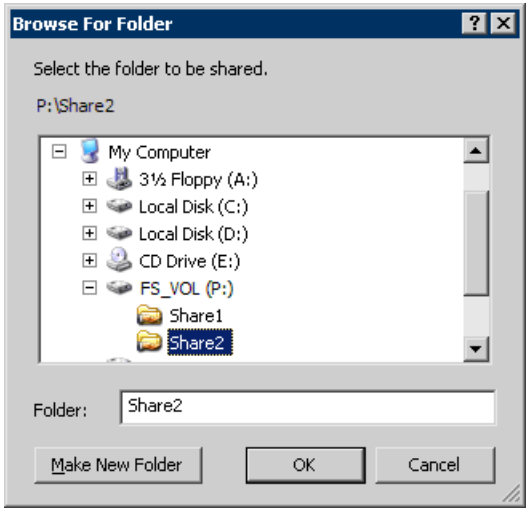

Virtual Computer Name	Type a unique virtual computer name to be assigned to the file share server. This is the name that the clients use to access the file shares.The virtual name must not exceed 15 characters.
Path	Click the field and either type the path of the directory to be shared or click the ellipsis button (**...**) to browse for a directory. The selected directories must meet the following conditions: ■ 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**.

Configure NetApp SnapMirror Resource(s)	This is applicable in case of VCS for Windows only. Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

Click **Next**.

6 On the Share Permissions panel, specify the users for the file shares and assign permissions to them as follows:

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 to add the specified user to the **Selected Users** list. By default, all selected users are given the 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**.

Click **Next**.

7 On the Share Properties panel, set the share properties for the file shares as follows:

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 users allowed on Windows.
Enable cache	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: ■ **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 the 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	To apply the specified share properties to multiple file shares simultaneously, do the following: **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**. **Note:** This option is not visible if you are configuring only one share in the service group.

Click **Next**.

8 This is applicable in case of VCS for Windows only.

On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring multipath I/O (MPIO) over Fibre Channel (FC), you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

9 On the Network Configuration panel, specify information related to your network as follows:

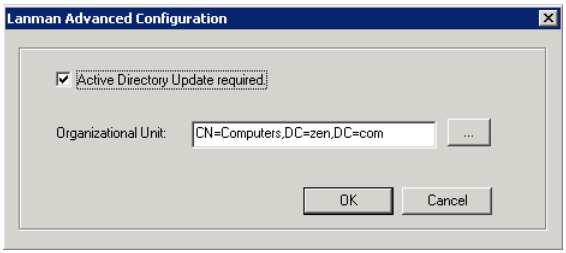

- Select **IPv4** to configure an IPv4 address for the virtual server.
  - In the **Virtual IP Address** field, type a unique virtual IPv4 address for the virtual server.

  - In the **Subnet Mask** field, type the subnet to which the virtual IPv4 address belongs.

- Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.
  - Select the prefix from the drop-down list. The wizard uses the prefix and automatically generates an IPv6 address that is valid and unique on the network.

- 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.

- Click **Advanced Settings** to specify additional details for the Lanman resource.

  On the Lanman Advanced Configuration dialog box, do the following:

  - 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.
  - 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 Organizational Unit, click the ellipsis button (**...**) and search using the Windows Find Organization Units dialog box.

    By default, the Lanman resource adds the virtual server to the default container "Computers."
  - Click **OK**.

    The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

Click **Next**.

10 On the Summary panel, review the service group configuration; the following service group details are displayed:

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.
Enable FastFailOver attribute for all the VMDg resources in the service group	This is applicable in case of SFW HA only. To enable all the VMDg resources in the service group for fast failover, select this checkbox.

Click **Next**.

11 Click **Yes** on the dialog that appears, informing you that the wizard will run commands to modify the service group configuration.

12 On the completion panel, 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**.

Creating non-scoped file shares configured with VCS

File shares configured with VCS on Windows Server are accessible only using the virtual server name (Lanman resource). These file shares are not accessible using the IP address.

The FileShare agent is enhanced to address this issue. The FileShare agent behavior can be controlled using the following registry key:

```
HKLM\SOFTWARE\VERITAS\VCS\BundledAgents\

Lanman\virtualName\DisableServerNameScoping
```

Set the DisableServerNameScoping key to have the FileShare agent support non-scoped file shares.

You must create this registry key manually.

Note: Incorrectly editing the registry may severely damage your system. Back up the registry before making changes.

To configure the DisableServerNameScoping registry key

1. To open the Registry Editor, press **Window+R** on the desktop (opens the Run dialog box), type `regedit`, and then click **OK**.
2. In the registry tree (on the left), navigate to the following location:

   ```
   HKLM\SOFTWARE\VERITAS\VCS\BundledAgents
   ```

3. Click **Edit > New > Key** and create a key by the name **Lanman**, if it does not exist already.
4. Select the **Lanman** key and click **Edit > New > Key** and create a key by the name ***virtualName***.

   Here, *virtualName* should be the virtual computer name assigned to the file share server. This is the VirtualName attribute of the Lanman resource in the file share service group.

   The newly created registry key should look like this:

   ```
   HKLM\SOFTWARE\VERITAS\VCS\BundledAgents\
   Lanman\virtualName
   ```

5. Select the key that you created in step 4 (***virtualName***) and add a DWORD type of value.

   The value name should be DisableServerNameScoping and value data should be 1.

   The value 1 indicates that the FileShare and Lanman agents support non-scoped file shares on Windows Server systems.
6. If there are multiple file share service groups to be used in the non-scoped mode, repeat steps 4 and 5 for each Lanman resource that is configured in the file share service group.
7. Save and exit the Registry Editor.

You must create this key only for Lanman resources that are part of VCS file share service groups. Configuring this key for Lanman resources that are part of other VCS service groups may result in unexpected behavior.

Making non-scoped file shares accessible while using virtual server name or IP address if NetBIOS and WINS are disabled

The VCS FileShare agent depends on NetBIOS or DNS to resolve the virtual name. If NetBIOS and WINS are disabled or the DNS is not updated, the agent is unable to resolve the virtual name.

This may typically occur when the file share service groups are configured to use localized IP addresses. When the service group is switched or failed over, the virtual name to IP address mapping changes. In such a case if WINS database or the DNS are not updated, the agent is unable to resolve the virtual name. As a result the FileShare resources fault and the shares become inaccessible.

The following message appears in the agent log:

```
VCS INFO V-16-10051-10530 FileShare:servicegroupname:online:
Failed to access the network path (\\virtualName)
```

The FileShare agent is enhanced to address this issue. The FileShare agent behavior can be controlled using the following registry key:

```
HKLM\SOFTWARE\VERITAS\VCS\BundledAgents\
\Lanman\virtualName\DisableStrictVirtualNameCheck
```

Set the DisableStrictVirtualNameCheck key to have the FileShare agent make the file shares accessible irrespective of whether or not the virtual name is resolvable. In case the virtual name is not resolvable, the file shares are accessible using the virtual IP.

You must create this registry key manually.

Note: Incorrectly editing the registry may severely damage your system. Back up the registry before making changes.

To configure the DisableStrictVirtualNameCheck registry key

1. To open the Registry Editor, press **Window+R** on the desktop (opens the Run dialog box), type `regedit`, and then click **OK**.
2. In the registry tree (on the left), navigate to the following location:

   ```
   HKLM\SOFTWARE\VERITAS\VCS\BundledAgents
   ```

3. Click **Edit > New > Key** and create a key by the name **Lanman**, if it does not exist already.

4 Select the **Lanman** key and click **Edit > New > Key** and create a key by the name ***virtualName***.

   Here, *virtualName* should be the virtual computer name assigned to the file share server. This is the VirtualName attribute of the Lanman resource in the file share service group.

   The newly created registry key should look like this:

   ```
   HKLM\SOFTWARE\VERITAS\VCS\BundledAgents\
   Lanman\virtualName
   ```

5 Select the key that you created in step 4 (***virtualName***) and add a DWORD type of value.

   The value name should be DisableStrictVirtualNameCheck and value data should be 1.

6 If there are multiple file share service groups to be used in the non-scoped mode, repeat steps 4 and 5 for each Lanman resource that is configured in the file share service group.

7 Save and exit the Registry Editor.

You must create this key only for Lanman resources that are part of VCS file share service groups. Configuring this key for Lanman resources that are part of other VCS service groups may result in unexpected behavior.

About configuring file shares with multiple subdirectories

Use the VCS CompositeFileShare agent to configure file shares with multiple subdirectories.

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

Before you configure multiple file shares

Note the following prerequisites before configuring file share using the CompositeFileShare agent:

- Verify that the directories to be shared are on shared disks or LUNs.
- Do not use local system accounts for share users. Use domain-level accounts and users only.

Note: Sharing a directory with a large number of subdirectories and enabling the ShareSubdirectories flag could cause increased failover time and high CPU and memory utilization.

Configuring multiple file shares

This topic describes how to configure multiple file shares:

To configure multiple files shares

1. Configure your shared storage.
2. Configure the NIC and IP resources.
3. Create a resource of type CompositeFileShare.
4. Configure the following required attributes for the CompositeFileShare resource:
   - MountResName: The name of the MountV resource on which the CompositeFileShare resource depends.
   - PathAndShareName: A list specifying the respective paths and share names of the directories to be shared. If the path of a shared directory is `\Documents`, and the share name is `UserDocs`, the attribute is defined in the configuration file as `PathandShareName is {"\\Documents" = "UserDocs"}`.

   To create a hidden share, set the HiddenShare attribute to 1. Do not append the share name with a $ (dollar) sign.

   For more information on other CompositeFileShare agent attributes, refer to the *Cluster Server Bundled Agents Reference Guide*.
5. Configure a Lanman resource. Do not create other resources on which the Lanman agent depends.
6. Link resources to create the following dependencies:
   - CompositeFileShare resource depends on the MountV (in case of SFW HA), Mount (in case of Windows LDM), or NetAppSnapDrive (in case of VCS for Windows) resource.
   - CompositeFileShare resources depends on the Lanman resource.
   - Lanman resource depends on IP or IPv6 resource, and the IP or IPv6 resource in turn depends on the NIC resource.
7. Configure other resources in the service group, if required.
8. Bring the Lanman resource, and other resources in the service group, online.

About configuring print shares

This topic provides an overview of the steps involved in configuring a print share service group in a VCS cluster. A print share service group enables clients to share a network printer from a cluster.

VCS provides several ways to configure a print share service group, including the configuration wizard, Cluster Manager (Java Console), and the command line. This section provides instructions on how to use the Print Share Configuration Wizard to configure print shares.

On Windows 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).

If UAC is enabled, run the program or commands in the “Run as administrator” mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

Before configuring a print share service group, review the resource types and attribute definitions of the PrintShare agents described in the *Cluster Server Bundled Agents Reference Guide*.

Note: The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems.

Before you configure a print share service group

Note the following prerequisites before you configure a print share service group:

- Verify that you have local administrator privileges on the node from where you run the wizard.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.
  For a detailed list of services and ports used, refer to the product installation and upgrade guide.
- Verify that the VCS high availability engine, HAD, is running on the node from which you run the wizard.
- Verify that VCS Command Server service is running on all systems in the cluster.
- Verify that the network printer has an IP address assigned.

- Veritas recommends creating the spooler and the replication directories on different disk partitions, volumes, or LUNs.
- Mount the drives or LUNs with the spooler and the replication directories on the node on which you run the wizard. Unmount the drives or LUNs from other nodes 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). VEA is available with SFW HA only.
- If NetBIOS is disabled over TCP/IP, you must set the Lanman agent's DNSUpdateRequired attribute value to 1 (True).
  You can modify the Lanman resource attribute value after configuring the service group.
- Verify that you have the following information ready. The wizard prompts you for these details:
  - A unique virtual computer name to be assigned to the print share server
    This is the name by which clients access the server. The virtual name must not exceed 15 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 virtual IP address is required only if you wish to configure an IPv4 address. In case of IPv6, the wizard prompts you to select the IPv6 network and automatically generates an IPv6 address that is valid and unique on the network. The wizard uses the prefix that is advertised by the router on the IPv6 network.
  - The IP address of the network printer
- To be able to add or modify custom settings on printers clustered under VCS, you must perform the following tasks manually:
  - Create the following DWORD registry key and set its value to `1`.
    For global configurations, create:

    ```
    HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\BundledAgents\
    PrintSpool\__GLOBAL__\PrinterDriverSupport
    ```

    For per-resource configurations, create:

    ```
    HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\BundledAgents\
    PrintSpool\<resource_name>\PrinterDriverSupport
    ```

  - After creating this key, do one of the following:

If a PrintSpool resource if offline, bring it online.

If a PrintSpool resource is already online, probe it.

The PrintSpool agent creates the`HKEY_LOCAL_MACHINE\Cluster` key. Make sure that the key exists, and then add printers to the virtual name.

By default, changes to registry keys are saved to the file system after every 5 monitor cycles. If you want to customize this setting, create the following DWORD registry key and set its value to a different number of monitor cycles:

For global configurations, create:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\BundledAgents\
PrintSpool\__GLOBAL__\PrinterDriverSupportFlushInterval
```

For per resource configurations, create:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\VCS\BundledAgents\
PrintSpool\<resource_name>\PrinterDriverSupportFlushInterval
```

Note: You might not be able to add or modify custom settings on existing printers clustered under VCS. In this case, Veritas recommends that you reinstall the printer drivers and add the Windows printer resources (not VCS printer resources) again.

Configuring a print share service group using the wizard

Configuring the print share service group involves creating a print share 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.

The Print Share Configuration wizard enables you to create and modify a print share service group in a VCS cluster. This section describes how to create a print share service group using the wizard.

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See “Configuring the service group in a non-shared storage environment” on page 163.

You can also modify an existing print share service group using the wizard.

See “Modifying a print share service group using the wizard” on page 168.

Configuring a print share service group involves the following tasks:

- Creating a service group with a PrintSpool resource and bringing it online. This also involves configuring the Lanman resource on which the PrintSpool resource depends.
- Adding a network printer to the virtual computer created by the Lanman resource, and creating a new TCP/IP port for the printer.
- Configuring a PrintShare resource in the service group and bringing it online.

These tasks are described in the procedures that follow.

To configure a print share service group with a PrintSpool resource

1. Start the Print Share Configuration Wizard.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center, expand **Solutions for Additional Applications**, and click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on 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**.

Specify the following details:

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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox.

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

Do the following:

- Select **IPv4** to configure an IPv4 address for the virtual server.
  - In the Virtual IP Address field, type a unique virtual IPv4 address for the virtual server.
  - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
- Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.
  Select a value from the **Prefix** from the drop-down list. The wizard uses this prefix and automatically generates an IPv6 address that is valid and unique on the network.

- In the Virtual Server Name field, type a unique virtual computer name by which the print share server will be known to clients. Note that the virtual name must not exceed 15 characters.
- 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.
- Click **Advanced Settings** to specify additional details for the Lanman resource.

  On the Lanman Advanced Configuration dialog box, do the following:

  - 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.
  - 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 **...** (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."
  - Click **OK**.

    The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

6 On the Configure Data Path panel, specify the directories for spool and registry replication, specify the other options on this panel, and then click **Next**.

Specify the following details:

Spooler Directory	Type the path or click **...** (ellipsis button) to browse for the directory. All print commands are 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 are 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. Veritas recommends creating the directories for replication and spooling on different mounts.
Configure NetApp SnapMirror Resource(s)	This is applicable in case of VCS for Windows only. Check the **Configure NetApp SnapMirror Resource(s)** check box if you wish to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

7 This step is applicable in case of VCS for Windows only.

On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring multipath I/O (MPIO) over Fibre Channel (FC), you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

8 On the Build Print Server panel, review the configuration and click **Next**.

The following service group details are visible:

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 Click **Yes** on the dialog that prompts you that the wizard will run commands to modify the service group configuration.

10 Bring the PrintSpool resource online.

Proceed to the next steps 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 the cluster.

   To launch the Add Printer wizard, return to the Print Share Configuration Wizard, and click **Add Printer** on the Add Printer panel.

   Alternatively, 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, and double-click **Add Printer**.

2 In the Add Printer wizard, review the information on 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**.

   Specify the following details:

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**.

The following service group details are visible:

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.
Enable FastFailOver attribute for all the VMDg resources in the service group	This is applicable in case of SFW HA only. To enable all the VMDg resources in the service group for fast failover, select this checkbox.

4 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**.

About configuring IIS sites

When configuring the IIS agent to monitor a Web site, you can monitor associated application pools in the following ways:

- Configure a resource to monitor the Web site and define options to monitor associated application pools within the same resource.
- Configure a resource to monitor the IIS site only. Configure additional resources to monitor specific application pools.

VCS provides several ways to configure the agent, including the configuration wizard, Cluster Manager (Java console), and the command line. This section provides instructions on how to use the wizard to configure sites.

To configure IIS agent on Windows Server Core, first install IIS 7.0 (for Windows Server 2008 R2 SP1) or IIS 8.0 (for Windows Server 2012 or Windows Server 2012 R2) in the specified order. Then, manually add the required resources and configure the service group.

See “Installing IIS 8.0 on Windows Server Core” on page 135.

See “Installing IIS 7.0 on Windows Server 2008 Server Core” on page 137.

Then, manually add the required resources and configure the service group.

See "Before you configure an IIS service group" on page 133.

See "Configuring an IIS service group using the wizard" on page 139.

You can perform the manual configuration steps either directly on the Server Core machine using VCS commands from the command line, or remotely using the Cluster Manager (Java console).

If UAC is enabled, run the program or commands in the "Run as administrator" mode even if the logged-on user belongs to the local administrators group. Alternatively, log on as an Administrator (default administrator account) to perform the tasks.

Before configuring the agent, review the agent's resource type definition and attribute descriptions in the *Cluster Server Bundled Agents Reference Guide*. Also, review the sample configurations and resource dependency graphs.

Before you configure an IIS service group

Note the following prerequisites before you configure an IIS service group:

- 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.
- For IIS 7.0 (on Windows Server 2008 R2 SP1) or IIS 8.0 (on Windows Server 2012 or Windows Server 2012 R2), you must install the following role services:
  - IIS 6 Metabase Compatibility
  - IIS 6 WMI Compatibility or the IIS Management Scripts and Tools
    Only one of these role services is required.

  These options are available under Management Tools on the Role Services page of the Add Roles Wizard.

  If IIS 6 Metabase Compatability role is installed, the WMI 6 Provider is used. If IIS Management Scripts and Tools role is installed, the WMI 7 Provider is used. If both the roles are installed, the WMI 7 Provider is used.

  These components are required for the IIS agent to function on Windows Server.

  For the Windows Server Core editions, you must install IIS 7.0 (on Windows Server 2008 R2 SP1) or IIS 8.0 (on Windows Server 2012 or Windows Server 2012 R2) in the specified order.

  See "Installing IIS 8.0 on Windows Server Core" on page 135.
  See "Installing IIS 7.0 on Windows Server 2008 Server Core" on page 137.
- If IIS configuration is using IPv6 addresses, then you must install the IIS Management Scripts and Tools role service.

IPv6 requires WMI 7 Provider that is part of the IIS Management Scripts and Tools role.

- If you are configuring FTP sites that use IPv6 addresses, ensure that the IPv6 address entry (IP Address column in Site Bindings dialog) is enclosed in square brackets. The VCS IIS Configuration Wizard requires this format to correctly configure the FTP site in the cluster.
  See "Fixing the IPv6 address configuration for FTP sites" on page 135.
- 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 that 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 "About configuring IIS sites" on page 132.
- Verify that you have local administrator privileges on the system from where you run the wizard.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.
  For a detailed list of services and ports used refer to the product installation and upgrade guide.
- Verify that the VCS engine, HAD, is running on the node from which you run the wizard.
- Mount the drives or LUNs containing the shared directories on the node from which you run the wizard. Unmount the drives or LUNs from other nodes 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). VEA is available with SFW HA only.
- Keep the following information ready. The wizard prompts you for these details:
  - 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.

Fixing the IPv6 address configuration for FTP sites

When you add an FTP site using the Add FTP Site wizard, the IPv6 address is not enclosed in brackets by default. The VCS IIS Configuration Wizard requires the IPv6 addresses enclosed in square brackets format to correctly configure the FTP site in the cluster.

1. From the IIS Manager, right-click the FTP site name and click **Bindings**.
2. In the Site Bindings dialog box, select the FTP site and click **Edit**.
3. In the Edit Site Binding dialog box, type square brackets around the IPv6 address displayed in the IP address field.

   For example, the IPv6 address should display as `[2001:Db8:0:10:828:1871:cd8:5c0f]`.
4. Click **OK** and then click **Close**.

Installing IIS 8.0 on Windows Server Core

On Windows Server Core, you must install IIS in the order specified in this procedure.

To install IIS 8.0 on Windows Server Core

1 Type the following at the command prompt:

```
C:\>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-Win dowsAuthentication;
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-FTPServer
```

2 Verify that all the components specified in the earlier step have successfully installed. Type the following at the command prompt:

```
C:\>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:

```
C:\>oclist
```

5 Repeat the steps on all the nodes where you want to configure the IIS service group.

Installing IIS 7.0 on Windows Server 2008 Server Core

On Windows Server 2008 Server Core, you must install IIS in the order specified in this procedure.

To install IIS 7.0 on Windows Server 2008 Server Core

1. Type the following at the command prompt:

```
C:\>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-Win dowsAuthentication;
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-FTPServer
```

2. Verify that all the components specified in the earlier step have successfully installed. Type the following at the command prompt:

```
C:\>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:

```
C:\>oclist
```

5 Repeat the steps on all the nodes where you want to configure the IIS service group.

Configuring an IIS service group using the wizard

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.

The IIS Configuration Wizard enables you to create and modify IIS service groups, making sites highly available in VCS cluster.

The wizard creates one resource for each IIS site and its associated application pools; the wizard does not create resources that monitor only application pools.

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See "Configuring the service group in a non-shared storage environment" on page 163.

To configure an IIS service group using the wizard

1 Start the IIS Configuration Wizard.

or

Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC). In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **IIS Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu.

2 Review the information on 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**.

Specify the following details:

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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox.

5 On the Configure IIS Sites panel, add and remove sites from the service group, configure IP addresses, ports, and virtual computer names, optionally choose to configure NetApp SnapMirror resources and then click **Next**.

Specify the following details:

Add	Check the check box corresponding to the site to be configured in VCS.
IP	Verify or 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.
Virtual Name	Type a virtual name for the selected site. Each virtual name can be associated with only one virtual IP address at a time.
Configure NetApp SnapMirror Resource(s)	This is applicable with VCS for Windows only. Check the **Configure NetApp SnapMirror Resource(s)** check box if you want to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.

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

Specify the following details:

IP Address	Displays the virtual IP addresses. The wizard groups systems by the virtual IP addresses associated with the systems.
Subnet Mask	If the virtual IP is an IPv4 address, verify or type the subnet mask associated with each virtual IPv4 address. If the virtual IP is an IPv6 address, verify or type the associated IPv6 prefix. The prefix is generally represented in the following format: `ipv6-address/prefix-length`. For example: `2001:db8:0:1::/64`
Adapter Name	Select the public adapter associated with the virtual IP address on each system.

7 This is applicable with VCS for Windows only.

On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring multiPath I/O (MPIO) over Fibre Channel (FC), you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

Specify the following details:

Site Name	Displays the site names.

AppPoolMon	For each site, select the monitoring options from the AppPoolMon list. Choose from the following options from the drop-down list: ■ **NONE**—The agent does 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.

9 On the Service Group Summary panel, review the service group configuration and then click **Next**.

The following service group details are visible:

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.
Enable FastFailOver attribute for all the VMDg resources in the service group	This is applicable to SFW HA only. To enable all the VMDg resources in the service group for fast failover, select this checkbox.

10 Click **Yes** on the dialog that prompts you that the wizard will run commands to modify the service group configuration.

11 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**.

About configuring applications using the Application Configuration Wizard

VCS provides an Application Configuration Wizard to create service groups to monitor applications that are configured as resources of type GenericService, ServiceMonitor, or Process. You can also use the wizard to add registry replication and network resources to application service groups.

Note: The wizard does not configure the registry replication and network resources independently. It configures these resources as part of a service group that has application resources.

On Windows 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 VCS commands, or remotely using the Cluster Manager (Java console).

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

Before you configure service groups using the Application Configuration wizard

Note the following prerequisites before you configure application service groups using the Application Configuration wizard:

- Verify that the application you wish to configure is installed on the nodes that are going to be part of the service group.
- Verify that the startup type of the application service that you wish to configure is set to manual on all the nodes that are going to be part of the service group.
- Verify that the application service is stopped on all the nodes that are going to be part of the service group.
- Verify that the shared drives or LUNs required by the applications are mounted on the node where you run the wizard.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.

  For a detailed list of services and ports used, refer to the product installation and upgrade guide.
- Before running the wizard, make sure you have the following information ready:
  - Details of the application that you wish to configure (for example, application type, service name, start parameters, startup directory)
  - Shared storage used by the applications
  - Application registry entries for configuring registry replication
  - Network and virtual computer (Lanman) details for the application

Note: These prerequisites apply to Application Configuration Wizard. For agent-specific prerequisites, see the agent descriptions in the *Cluster Server Bundled Agents Reference Guide*.

Adding resources to a service group

This topic describes how to use the Application Configuration Wizard to add resources to a service group.

To add resources to a service group

1. Start the Application Configuration Wizard.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC). In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Application Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on 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 following service group details and then click **Next**:

Service Group Name	Type a name for the 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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox.

5 The Application Options dialog box provides you the option to specify the type of application to be configured.

The following options are available:

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 147.
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 148.
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 “Adding resources to a service group” on page 145.

Configuring a GenericService resource

This topic descibes how to use the Application Configuration Wizard to configure a GenericService resource.

To configure a GenericService resource

1. In the Application Options panel, click **Create**, select **GenericService** from the corresponding drop-down list, and click **Next**.
2. On the Generic Service Options panel, specify the details of the service that you wish to configure and then click **Next**.

   Specify the service for which you wish to configure a GenericService resource and then specify the following attributes:

   - Click the **...** (ellipsis button) 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 function.
   - 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 function.
3. On the User Details panel, specify the details of the user in whose context the service will run and then click **Next**.

   Do the following:

   - 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** and then specify the **Domain Name**, **User Name**, and **Password** in the respective fields.
4. On the Shared Storage Option panel, under Available Shared Drives box, select the check box adjacent to the shared drive and then click **Next**.

   This is the shared storage that is required by the GenericService resource. The shared storage that you select will be in addition to the mount where the service binaries exist.
5. In the Application Resource Summary panel, review the summary of the GenericService resource. Click **Back** to make changes. Otherwise, click **Next**.
6. In the Application Options dialog box, select one of the following options:

- To configure another GenericService resource, repeat step To configure a GenericService resource through step To configure a GenericService resource.
- To configure a Process resource:
  See "Configuring processes" on page 148.
- To configure a ServiceMonitor resource:
  See "Configuring a ServiceMonitor resource" on page 150.
- To configure other resources, including FileShare, Registry Replication, and Network resources:
  See "Configuring VCS components" on page 151.

If you do not wish to add any more resources, proceed to configuring the service group.

See "Configuring service groups using the Application Configuration Wizard" on page 154.

Configuring processes

This topic describes how to use the Application Configuration Wizard to configure processes.

To configure processes

1. In the Application Options panel, click **Create**, select **Process** from the corresponding list, and click **Next**.
2. On the Process Details panel, specify the details of the process that you wish to configure and then click **Next**.

   Specify the process details as follows:

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

- In the Monitor Program text box, type the complete path of the program that monitors the Start Program or browse for it by clicking **...** (ellipsis button).
  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, type the complete path of the Clean process or browse for it by clicking **...** (ellipsis button).
  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.
- Check 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.

3 On the User Details panel, specify information about the user in whose context the process will run and then click **Next**.

Do the following:

- 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** and then specify the **Domain Name**, **User Name**, and **Password** in the respective fields.
- Click **Next**.

4 On the Shared Storage Option panel, under Available Shared Drives box, select the check box adjacent to the shared drive and then click **Next**.

This is the shared storage required by the Process resource. The shared storage that you select will be in addition to the mount where the process binaries exist.

5 In the Application Resource Summary panel, review the summary of the Process resource. Click **Back** to make changes. Otherwise, click **Next**.

6 In the Application Options dialog box, 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 147.

- To configure a ServiceMonitor resource:
  See "Configuring a ServiceMonitor resource" on page 150.
- To configure other resources, including Registry Replication and Network resources:
  See "Configuring VCS components" on page 151.
  If you do not want to add any more resources, proceed to configuring the service group.
  See "Configuring service groups using the Application Configuration Wizard" on page 154.

Configuring a ServiceMonitor resource

This topic descibes how to use the Application Configuration Wizard to configure a ServiceMonitor resource.

To configure 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, do the following:

   - Select the **Service** option and click **...** (ellipsis button) 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 the service name to be monitored.
   - Click **Next**.

   If you want a script to monitor the service, do the following:

   - Click **...** (ellipsis button) and specify the complete path for the script.
   - 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. On the User Details panel, specify the user information in whose context the service will be monitored.

   Do the following:

   - 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** and then specify the **Domain Name**, **User Name**, and **Password** for the user account.

  If the service selected in step 2 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**.

  Service Monitor 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 dialog box 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 In the Application Options dialog box, 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 147.

- To configure a Process resource:

  See "Configuring processes" on page 148.

- To configure other resources, including Registry Replication and Network resources:

  See "Configuring VCS components" on page 151.

  If you do not want to add any more resources, proceed to configuring the service group.

  See "Configuring service groups using the Application Configuration Wizard" on page 154.

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 as follows:

   - Registry Replication Component: Select this option to configure registry replication for your application. To configure a Registry Replication resource, proceed to step 3.
   - 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 5.

   The wizard does not enable the **Lanman Component** check box unless the **Network Component** check box is checked.
3. Specify the registry keys to be replicated.

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

   - 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.
   - This is applicable in case of VCS for Windows only.
     Check the **Configure NetApp SnapMirror Resource(s)** check box if you want to set up a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site.
   - Click **Next**.

   If you chose Network Component from the Application Component dialog box, proceed to the next step. Otherwise, proceed to step 6.

4 This step is applicable in case of VCS for Windows only.

On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring multipath I/O (MPIO) over Fibre Channel (FC), you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

Specify the network related information as follows:

- Select **IPv4** to configure an IPv4 address for the virtual server.
  - In the Virtual IP Address field, type a unique virtual IPv4 address for the virtual server.
  - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
- Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.
  - Select the prefix from the drop-down list. The wizard uses the prefix and automatically generates an IPv6 address that is valid and unique on the network.
- In the Virtual Server Name field, enter a unique virtual computer name by which the node will be visible to the other nodes.

  The virtual name must not exceed 15 characters. Note that the Virtual Computer Name text box is displayed only if you chose to configure the Lanman Component in Application Component dialog box.
- 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. Ensure that you select the adapters assigned to the public network, not the private.
- Click **Advanced** and then specify additional details for the Lanman resource as follows:
  - 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 **...** (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."

  The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.

- Click **Next**.

6 In the Application Options dialog box, select one of the following options:

- To configure additional VCS components, repeat step 1 through step 5.
- To configure a GenericService resource:
  See "Configuring a GenericService resource" on page 147.
- To configure a Process resource:
  See "Configuring processes" on page 148.
- To configure a Service Monitor resource:
  See "Configuring a ServiceMonitor resource" on page 150.

If you do not want to add any more resources, proceed to configuring the service group:

See "Configuring service groups using the Application Configuration Wizard" on page 154.

Configuring service groups using the Application Configuration Wizard

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.

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

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See "Configuring the service group in a non-shared storage environment" on page 163.

To configure a service group using the wizard

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

   The option is enabled only if the following conditions are met:

   - 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 this dialog box 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 dialog box 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 On the Service Group Summary panel, review the service group configuration and click **Next**.

The following service group details are visible:

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.
Enable FastFailOver attribute for all the VMDg resources in the service group	This is applicable in case of SFW HA only. To enable all the VMDg resources in the service group for fast failover, select this checkbox.

4 Click **Yes** on the dialog that prompts you that the wizard will run commands to modify the service group configuration.

5 In the completion panel, check **Bring the service group online** if you want to bring the service group online on the local system.

6 Click **Finish** to create the service group and exit the Application Configuration Wizard.

About configuring the Oracle service group using the wizard

Configuring the Oracle database agent involves creating the Oracle service group, its resources, and defining attribute values for the configured resources.

VCS provides a configuration wizard that guides you through the process of configuring the Oracle service group. You can use the wizard to create and modify Oracle service groups.

This section describes the steps required to create a new Oracle service group using the wizard.

Review the instructions to modify an existing service group.

See "About modifying an Oracle service group" on page 172.

Prerequisites for configuring the Oracle service group

Prerequisites for configuring the Oracle service group are as follows:

- If you have configured a firewall, add the required services and ports to the Firewall Exception list.
  For a detailed list of the required services and ports, see the *Veritas InfoScale Installation and Upgrade Guide*.
- Verify that the appropriate InfoScale product is installed on all cluster nodes.
- Verify a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
  See "Configuring the cluster using the Cluster Configuration Wizard" on page 345.
- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be logged on as a Domain Administrator on the node where you run the wizard.
- Verify that the Veritas High Availability Engine (HAD) is running on the system from where you run the wizard.
- Connect the LUNs or mount the volumes containing the data files, control files, redo log files, bdump, cdump, and udump files. Unmount or disconnect the volumes or LUNs from other nodes in the cluster.
- Mount the database and start the Oracle instance on the node running the wizard.
- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved in the service group configuration.
  After you modify the service group configuration you can again edit the resource attributes to set the desired value.
- Keep 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

- 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.

Creating an Oracle service group using the wizard

This section describes how to create an Oracle service group.

If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

See "Configuring the service group in a non-shared storage environment" on page 163.

To create an Oracle service group

1. Start the Oracle configuration wizard.

   In case of VCS, Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu.

   In case of SFW HA, click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu to start the Solutions Configuration Center (SCC). In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.
2. On the Welcome panel click **Next**.
3. On the Wizard Options panel, select **Create service group** and click **Next**.
4. On the Service Group Configuration panel, use the following fields as appropriate:

Service Group Name	Type a name for the Oracle service group.

Group System List	In the Available Cluster Systems list, 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 list. 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.
Include selected systems in the service group's AutoStartList attribute	To enable the service group to automatically come online on one of the systems, select this checkbox. For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.

Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.

5 On the Oracle Configuration panel, complete the following and then click **Next**.

- Select the SIDs and the associated listeners to be added to the service group. 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**.

Enter the appropriate information in the following fields:

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\`. Here, *%vcs_home%* is the installation directory, typically `C:\Program Files\Veritas\Cluster Server`.

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.

Do the following:

- In case of IPv4, select **IPV4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- In case of IPv6, select **IPV6** and select the IPv6 network from the drop-down list.

  The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.

  The IPv6 option is disabled if the network does not support IPv6.
- For each system in the cluster, select the public network adapter name. The Adapter Display Name 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.

Note: If you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that have the same MAC address, then you must edit the “MACAddress” attribute of the NIC agent and the IP agent, after you configure the application service group.

---

9 Review the configuration on the Summary panel.

Resources	Lists the configured resources. The wizard assigns unique names to the resources. Click on a resource to view its attributes and their configured values in the Attributes box.
Attributes	Enables you to edit a resource name. Click the resource name or press the F2 key. After editing, press the Esc key to cancel the changes, or press the Enter key to confirm the changes.
Enable FastFailOver attribute for all the VMDg resources in the service group	To enable all the VMDg resources in the service group for fast failover, select this checkbox.

Click **Next**.

10 On 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 *Cluster Server Administrator's Guide* for configuration instructions.

Enabling fast failover for disk groups (optional)

For service groups that contain many disk groups, you can greatly reduce failover time by implementing the SFW fast failover feature for disk groups.

More information is available about fast failover benefits and requirements.

See "Considerations for a fast failover configuration" on page 73.

For implementing the fast failover feature, VCS provides a new attribute, FastFailOver, for the Volume Manager Diskgroup (VMDg) resource. This attribute determines whether or not a disk group is enabled for fast failover.

Note: The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click **Properties**. Disk group version upgrade is required after upgrading SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

You can enable fast failover for all the VMDg resources while configuring the service group using the configuration wizard. The service group configuration wizard provides a checkbox to enable fast failover.

Perform these steps if you did not enable fast failover using the wizard or if you have configured the service group manually.

The following procedure describes how to enable the FastFailOver attribute using the VCS Java Console.

To enable the FastFailover attribute for a VMDg resource

1. In Cluster Manager (Java Console), select a service group with a VMDg resource configured for it.

   Select the Properties tab from the right pane.
2. Scroll down to choose the **FastFailOver** attribute and click to edit the attribute value.

3 In the Edit Attribute dialog box, check the **FastFailOver** check box and then click **OK**.

4 Repeat these steps for every VMDg resource for which you want to enable fast failover.

Configuring the service group in a non-shared storage environment

If you are using a non-shared storage configuration, you have to use the VCS MountV – VMNSDg agents to monitor your local storage. Currently, the service group configuration wizards do not support configuring these agents in the service group. You have to configure these agents manually by using the Cluster Manager (Java Console) or the VCS commands.

VCS provides templates for configuring service groups that use non-shared storage agent resources.

The Java Console templates are located in the following directory:

```
%VCS_HOME%\Templates
```

Here, *%VCS_HOME%* is the default product installation directory, typically, `C:\Program Files\Veritas\Cluster Server`.

For information about adding a service group using templates from the Java Console, refer to the *Cluster Server Administrator's Guide*.

The following steps describe how to create a service group using the Cluster Manager (Java Console).

To configure the service group in a non-shared storage environment

1 Open the Cluster Manager (Java Console) from **Start > All Programs > Veritas > Veritas Cluster Server** and then click **Veritas Cluster Manager - Java Console** or, on Windows Server 2012 operating systems, from the **Apps** menu.

2 Log on to the cluster. On the Cluster Monitor window click **File > New Cluster**, then on the New Cluster window type **localhost** in the Host name field, and then click **OK**.

3 Launch the service group configuration wizard. From the Cluster Explorer window menu, click **Tools > Configuration Wizard**.

4 On the Service Group Configuration Wizard Welcome panel, click **Next**.

5 Fill in the following information and then click **Next**:

- Specify a name for the service group.

- Select the systems for the service group. Click a system in the Available Systems box and then click the right arrow to move the systems to Systems for Service Group.
- Leave the service group type as the default, Failover.

6 Click **Next** again.

7 In the Templates list, select the desired service group template depending on the configuration and then click **Next**.

Template name	Description
FileShareVMNSGroup **PrintShareVMNSGroup** **IIS60VMNSGroup** **MSMQVMNSGroup**	Use these templates to create a single node high availability service group that uses non-shared storage. These templates include resources for configuring MountV and VMNSDg agents.
FileShareVirtVMNSGroup **PrintShareVirtVMNSGroup** **IIS60VirtVMNSGroup** **MSMQVirtVMNSGroup**	Use these templates to create a single node high availability service group in a VMware virtual environment. These templates includes resources for configuring MountV, VMwareDisks, and VMNSDg agents.
VvrRvgVMNSRVGGroup	Use this template to create a Volume Replicator replication service group on a single node that uses non-shared storage.

The Templates box lists the templates available on the system to which Cluster Manager is connected. The resource dependency graph of the templates, the number of resources, and the resource types are also displayed.

8 Click **Next**. The wizard starts creating the service group.

9 After the service group is successfully created, click **Next** to edit attributes using the wizard.

10 The wizard lists the resources and their attributes. You must specify values for the mandatory attributes that appear in bold. The remaining resources listed in the window are preconfigured by the template and do not require editing.

To modify an attribute, do the following:

- Click the resource.
- Click the attribute to be modified.
- Click the **Edit** icon at the end of the table row.

- In the Edit Attribute dialog box, enter the attribute values.
- Click **OK**.

For details on application-specific agent attributes, refer to the application-specific agent or solutions guide.

For details on the storage and network agent attributes, refer to the *Cluster Server Bundled Agents Reference Guide*.

11 Click **Finish**.

12 Right-click the newly created service group and select **Enable Resources**.

13 Right-click the newly created service group, select **Online** from the context menu, and then select a system on which to bring the service group online.

If you are configuring the service group on a node at the secondary site in a DR environment, bring the service group online only after completing all the DR configuration steps.

Verifying the cluster configuration

Simulating a failover is an important part of configuration testing. 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, as follows:

- 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 that you selected to switch to in the first step.

3 To move all the resources back to the original node, repeat the first step of this procedure 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, perform these steps sequentially:

- Restart the node that you shut down in the first step.
- 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:

- Modifying the existing cluster configuration to add additional nodes or remove existing nodes.
  See "Adding nodes to a cluster" on page 271.
- Modifying the application service group configuration.
  See "Modifying the application service groups" on page 166.

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:

- See "Modifying a file share service group using the wizard" on page 167.
- See "Modifying a print share service group using the wizard" on page 168.
- See "Modifying an IIS service group using the wizard" on page 169.
- See "Modifying an application service group" on page 170.
- See "About modifying an Oracle service group" on page 172.

Modifying a file share service group using the wizard

The File Share Configuration Wizard enables you to modify a file share service group.

Consider the following before you modify file share service groups using the wizard:

- If the file share service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change attributes of resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), Mount and DiskRes (in case of Windows LDM), and NetAppSnapDrive and NetAppFiler (in case of VCS for Windows) 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.
- After configuring a file share if you move the shared directory to a new location, then while reconfiguring the file share service group, the wizard fails to delete the storage resources configured for the existing file share.

  The wizard successfully creates a new file share resource and the corresponding storage resources, but fails to remove the older storage resources from the service group.

  In such cases, you can either remove the stale storage resources manually, or delete the file share service group and run the wizard again to recreate the service group.

To modify a file share service group using the wizard

1. Start the File Share Configuration Wizard on a node on which the file share service group is online.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center, expand **Solutions for Additional Applications** and click **High Availability (HA) Configuration > Configure the Service Group > File Share Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on the Welcome panel and click **Next**.
3. On the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.
4. Follow the wizard instructions and make desired modifications to the service group configuration.

   See "About configuring file shares" on page 110.

Modifying a print share service group using the wizard

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

Consider the following before you modify print share service groups using the wizard:

- If the print share service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change attributes of resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), Mount and DiskRes (in case of Windows LDM), and NetAppSnapDrive and NetAppFiler (in case of VCS for Windows) 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 using the wizard

1. Start the Print Share Configuration Wizard on a node on which the print share service group is online.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center, expand **Solutions for Additional Applications**, and click **High Availability (HA) Configuration > Configure the Service Group > Print Share Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Read the information on the Welcome panel and click **Next**.

3. On the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.

4. Follow the wizard instructions and make desired modifications to the service group configuration.

   See "Configuring a print share service group using the wizard" on page 126.

   If you are modifying the service group to remove a PrintShare resource, make sure you offline the resource before deleting it.

Modifying an IIS service group using the wizard

The IIS Configuration Wizard enables you to modify an IIS service group.

Consider the following before you modify an IIS service group:

- If the IIS service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change attributes of resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), Mount and DiskRes (in case of Windows LDM), and NetAppSnapDrive and NetAppFiler (in case of VCS for Windows) 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.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**. From the Solutions Configurations Center, expand **Solutions for Additional Applications** and click **High Availability (HA) Configuration > Configure the Service Group > IIS Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on the Welcome panel and click **Next**.

3. On the Wizard Options panel, click **Modify service group**, select the service group to be modified, and click **Next**.

4. Follow the wizard instructions and make the modifications that you want to the service group configuration.

   See "Configuring an IIS service group using the wizard" on page 139.

Modifying an application service group

You can modify an application service group using the Application Configuration Wizard.

Consider the following before you modify service groups using the wizard:

- 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 use the wizard to add or remove resources from the configuration. You cannot modify resources that are online.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg (in case of SFW HA), Mount and DiskRes (in case of Windows LDM), and NetAppSnapDrive and NetAppFiler (in case of VCS for Windows) 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: Veritas recommends that you do not use the wizard to modify service groups that were not created using the wizard.

To modify an application service group

1. Start the Application Configuration Wizard.

   or

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC). In the SCC, open the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Application Configuration Wizard**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Review the information on the Welcome panel and click **Next**.

3. On 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 and click **Next**.

4. On the Service Group Configuration panel, if required, make changes as appropriate to update the SystemList and AutoStartList attributes, and then click **Next**.

   If you want the service group to automatically come online on one of the systems, make sure to select the **Include selected systems in the service group's AutoStartList attribute** checkbox.

5. Click **Modify**, select the resource you want to modify and then click **Next**.

   The Modify option is enabled only if the following conditions are met:

   - Service and Process resources are already configured using the wizard.
   - You selected the **Modify Service Groups** option in the Wizard Options panel.

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 the Service Monitor Options dialog box.

   Make required changes in the appropriate dialog box and click **Next**.

   See "Configuring a GenericService resource" on page 147.

   See "Configuring processes" on page 148.

   See "Configuring a ServiceMonitor resource" on page 150.

7. On the User Details dialog box, specify the user information and click **Next**.

8 On the Application Resource Summary dialog box, review the summary of the resource.

When modifying a volume in the service group, the **Enable FastFailOver attribute for all the VMDg resources in the service group** checkbox appears. Select this checkbox to enable all the VMDg resources in the service group for fast failover.

Click **Back** to make changes. Otherwise, click **Next**.

9 Repeat step 5 through step 8 for each resource that you want to modify.

10 After modifying the required resources, you can:

- Add additional resources to the service group.
  See "Adding resources to a service group" on page 145.
- Delete resources from the service group.
- Add VCS components to the service group.
  See "Configuring VCS components" on page 151.
- Create the service group.
  See "Configuring service groups using the Application Configuration Wizard" on page 154.

About modifying an Oracle service group

The section below describe how to modify the configuration of the service groups using the configuration wizard.

Prerequisites for modifying the Oracle service group

Note the following prerequisites before modifying the Oracle service group:

- 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 use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
  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.
- 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.
- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved

in the service group configuration. After you modify the service group configuration you can again edit the resource attributes to set the desired value.

- If you have upgraded the product from the previous version to 7.1, then do not use the modify mode. Delete the service group and re-configure it using the Oracle configuration wizard.

Modifying an Oracle service group

To modify an Oracle service group

1. Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   or

   If you are in the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.

2. Review the prerequisites and click **Next**.

3. On 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.

   See "About configuring the Oracle service group using the wizard " on page 156.

Deleting an Oracle service group

The section below describes how to delete an Oracle service group using the configuration wizard.

To delete an Oracle service group

1. Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   or

   Launch the Solutions Configuration Center (SCC) from **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.
2. Review the prerequisites and click **Next**.
3. On the Wizard Options panel, click **Delete service group**, select the service group to delete, and click **Next**.
4. On the Service Group Summary panel, click **Next**.
5. On the dialog box 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 False.

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.

Veritas recommends that before configuring detail monitoring, you must ensure that Oracle is configured correctly and you are able to connect to the database.

Enabling detail monitoring

The section below describes how to enable detail monitoring using the configuration wizard.

To enable detail monitoring

1. Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   or

   Launch the Solutions Configuration Center (SCC) from **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen. In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.

2. Review the prerequisites and click **Next**.

3. On 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. On the Oracle Configuration panel, select the SIDs to be configured along with their respective listeners and click **Next**.

5. On the Detail Monitoring dialog box, specify information to enable detail monitoring, and click **Next**.

   Enter the appropriate information in the following fields:

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. On 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.

   See "About configuring the Oracle service group using the wizard " on page 156.

Disabling detail monitoring

The section below describes how to disable detail monitoring using the Oracle Configuration Wizard.

To disable detail monitoring

1. Start the Oracle configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Oracle Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   or

   Start the Solutions Configuration Center (SCC) from **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen. In the SCC, click the **Solutions** tab, expand **High Availability Configuration Wizards**, and click the **Launch** button for the **Oracle Agent Configuration Wizard**.

2. Review the prerequisites and click **Next**.

3. On 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. On the Oracle Configuration panel, click **Next**.

5. On 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.

   See "About configuring the Oracle service group using the wizard " on page 156.

Chapter 8

Adding DMP to a clustering configuration

This chapter includes the following topics:

- About Dynamic Multi-Pathing
- Overview of configuration tasks for adding DMP DSMs
- Reviewing the prerequisites
- Reviewing the configuration
- Setting up DMP in a new cluster configuration
- Setting up DMP in an existing cluster configuration

About Dynamic Multi-Pathing

Dynamic Multi-Pathing (DMP) is an optional software component in InfoScale products that provides redundant path support for your storage. This support is provided by DMP device-specific modules (DSMs). You can add DMP to a clustering configuration.

The steps for adding DMP 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 DMP 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 DMP 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 DMP software will automatically transfer the storage to the second path.

During the product installation, if you plan to install the DMP DSMs, you must connect only one host bus adapter path.

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 DSMs installed and running, 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 InfoScale Enterprise on the inactive node.
  Product installation 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 the prerequisites

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

For installation and configuration information, refer to the *Storage Foundation Administrator's Guide* and *Veritas InfoScale Installation and Upgrade Guide*.

Reviewing the configuration

The purpose of the configuration is to add Dynamic Multi-Pathing DSMs to a clustering configuration. DMP is an optional software component in InfoScale Enterprise, InfoScale Foundation, and InfoScale Storage that provides redundant path support for your storage. You need to have two switches per server, one switch to accommodate each path.

The following figure illustrates the four different host bus adapter paths.

Figure 8-1 Adding fault tolerance with DMP

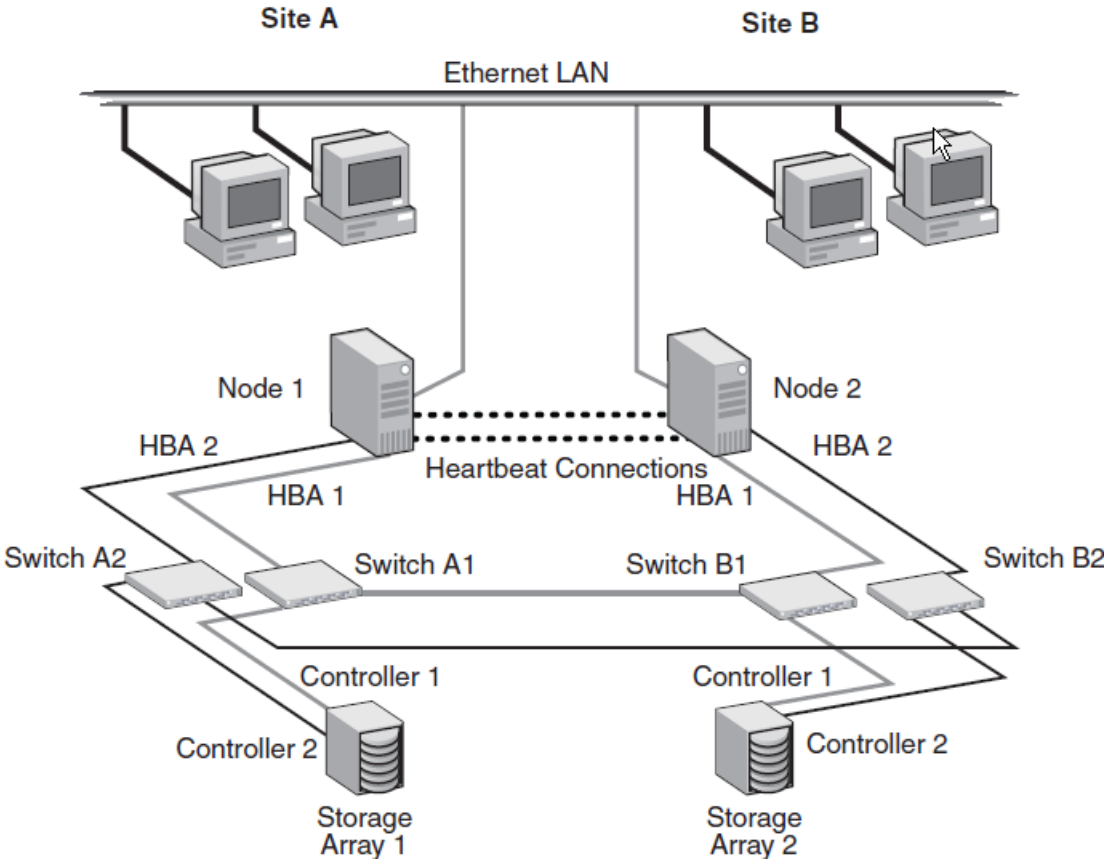

Setting up DMP in a new cluster configuration

Use the following procedure.

To set up DMP DSMs in a new cluster configuration

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 InfoScale Installation and Upgrade Guide*.
4. Reconnect the additional physical path.
5. Reboot the system.
6. Verify that the additional path is displayed.

Setting up DMP in an existing cluster configuration

Use the following procedure.

To set up DMP DSMs in an existing cluster configuration

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 DSMs using **Add or Remove Programs** from the Windows Control Panel.

   For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.
5. Reconnect the additional physical path.
6. Reboot the system.
7. Verify that the additional path is displayed.

Section 4

Campus Clustering

- Chapter 9. Introduction to campus clustering
- Chapter 10. Deploying InfoScale Enterprise for campus cluster

Chapter 9

Introduction to campus clustering

This chapter includes the following topics:

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

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

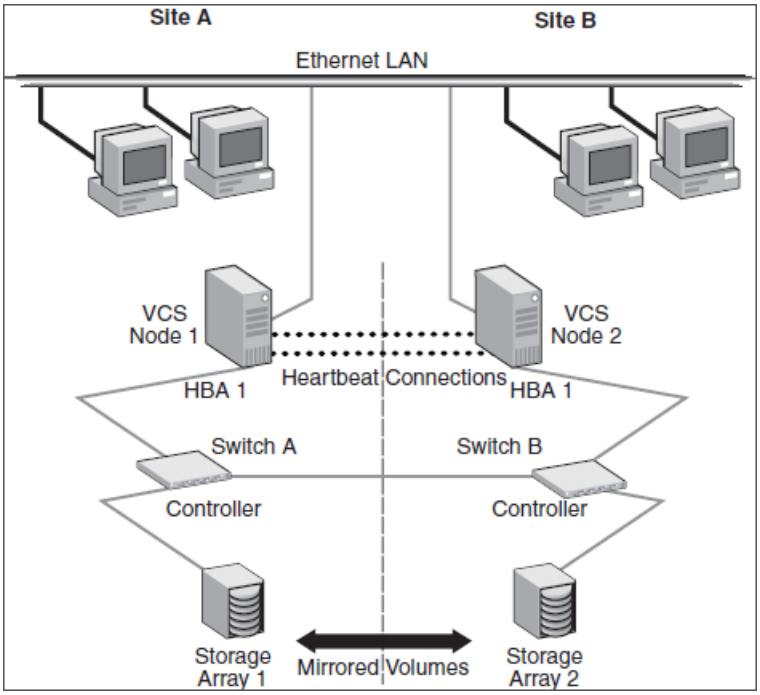

With SFW, a campus cluster can be set up using a Cluster Server (VCS) configuration. Both configurations involve setting up a single cluster with two nodes that are in separate buildings and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number 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.

Chapter 10

Deploying InfoScale Enterprise for campus cluster

This chapter includes the following topics:

- About the Campus Cluster solution
- Notes and recommendations for cluster and application configuration
- Reviewing the configuration
- Installing and configuring the hardware
- Configuring the storage hardware and network
- About installing the Veritas InfoScale products
- Configuring the cluster using the Cluster Configuration Wizard
- Creating disk groups and volumes
- Installing the application on cluster nodes
- Configuring service groups
- Verifying the cluster configuration

About the Campus Cluster solution

This chapter presents a VCS campus clustering configuration example with InfoScale Enterprise.

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

Table 10-1 Task list: Campus Cluster configuration

Objectives	Tasks
See "Reviewing the configuration" on page 190.	■ Review the configuration requirements ■ Overview of VCS campus cluster, and recovery scenarios
See "Installing and configuring the hardware" on page 195.	■ 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
See "Configuring the storage hardware and network" on page 335.	■ Install the operating system on both nodes ■ Make necessary networking settings on both nodes
See "About installing the Veritas InfoScale products" on page 503.	Install InfoScale Enterprise. Refer to the *Veritas InfoScale Installation and Upgrade Guide*.
See "Configuring the cluster using the Cluster Configuration Wizard" on page 345.	Use the VCS Cluster Configuration Wizard (VCW) to set up the cluster
See "Creating disk groups and volumes" on page 210.	■ Create dynamic cluster disk groups ■ Create dynamic volumes
See "Installing the application on cluster nodes" on page 220.	■ 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
See "Configuring service groups" on page 224.	■ Use an appropriate method to create and configure the VCS service group or groups ■ Bring the service group online
See "Verifying the cluster configuration" on page 269.	■ Switch the service group to the second node ■ Switch it back to the first node

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
  https://sort.veritas.com/documents

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

Refer to the following Microsoft knowledge database article for more details:

http://technet.microsoft.com/en-us/library/dd184075.aspx

- 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 *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. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

  Static IP addresses are required for the following purposes:

  - One static IP address 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 Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
  - For Volume Replicator 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 SFW HA 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 *Cluster Server Bundled Agents Reference Guide*.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- 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.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can

set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.

- For a Replicated Data Cluster, install only in a single domain.
- 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)
- This is applicable for a Replicated Data Cluster configuration.
  This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, 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.
- To configure a RDC cluster, you need to 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.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.

VCS agents, wizards, and other components	VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above). The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

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 192.

Figure 10-1 VCS campus clustering configuration example

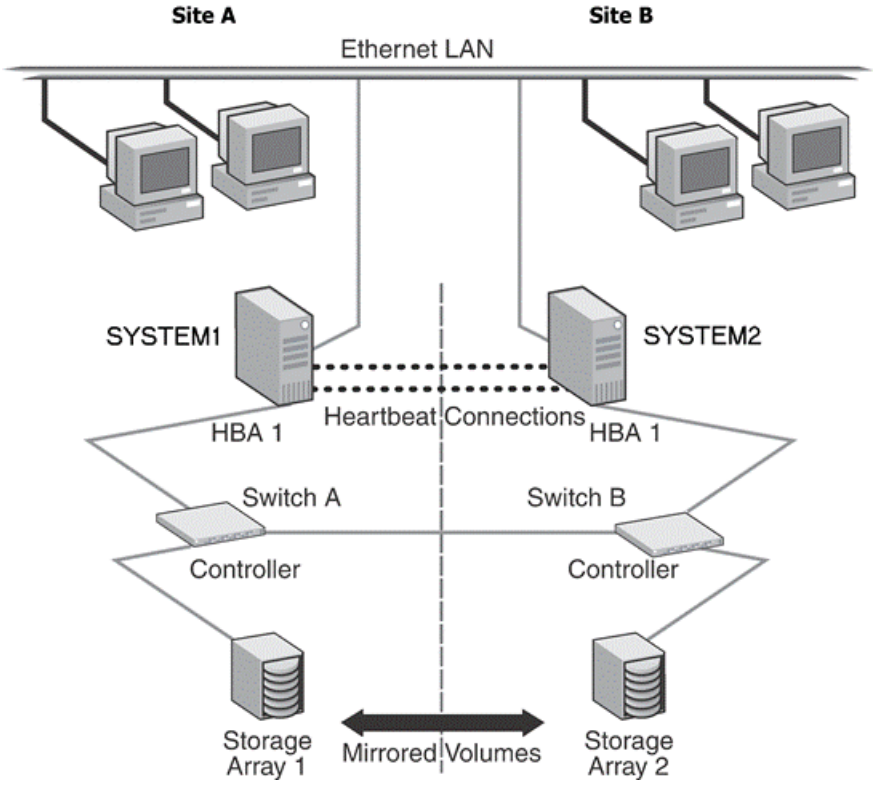

The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

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 "Overview of configuration tasks for adding DMP DSMs" on page 178.

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 following table 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).

For information on how to set the ForceImport attribute:

See "Setting the ForceImport attribute" on page 194.

Table 10-2 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 10-2 Failure situations *(continued)*

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
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.	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.
6. Storage interconnect lost Fibre interconnect severed.	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.
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.	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.

Refer to the failure scenarios:

See "Overview of campus clustering with VCS" on page 192.

For the failure scenarios 3 through 7 in this section, the following table 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 previous table.

Table 10-3 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 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. 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 vmdgResourceName 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 value 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
```

Here, 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, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Veritas 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.

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, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
- Veritas 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.

To verify the DNS settings and binding order for all systems

1. Open the Control Panel by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.

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, click **Adapter settings**.

4. Ensure the public network adapter is the first bound adapter by following these steps sequentially:

   - In the Network Connections window, click **Advanced > 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)** or **Internet Protocol Version 6 (TCP/IPv6)** check box, depending on which protocol your network is using.
   - 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 that the **Register this connection's address in DNS** check box is selected.

12 Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.

13 Click **OK**.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Configuring the cluster using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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 VCW to remove the node from the cluster, rename the system, and then run VCW again to add that system to the cluster.

Note the following prerequisites before you proceed:

- The required network adapters, and SCSI controllers are installed and connected to each system.

  To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.

- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.

- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

  Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list.
  When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
  Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.

- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

1. Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
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.

   To discover information about all systems and users in the domain, do the following:

   - Clear **Specify systems and users manually**.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - Select **Specify systems and users manually**.
     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. 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.

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**.

   Select the 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.

   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.
   - Product is either not installed or there is a version mismatch.

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

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

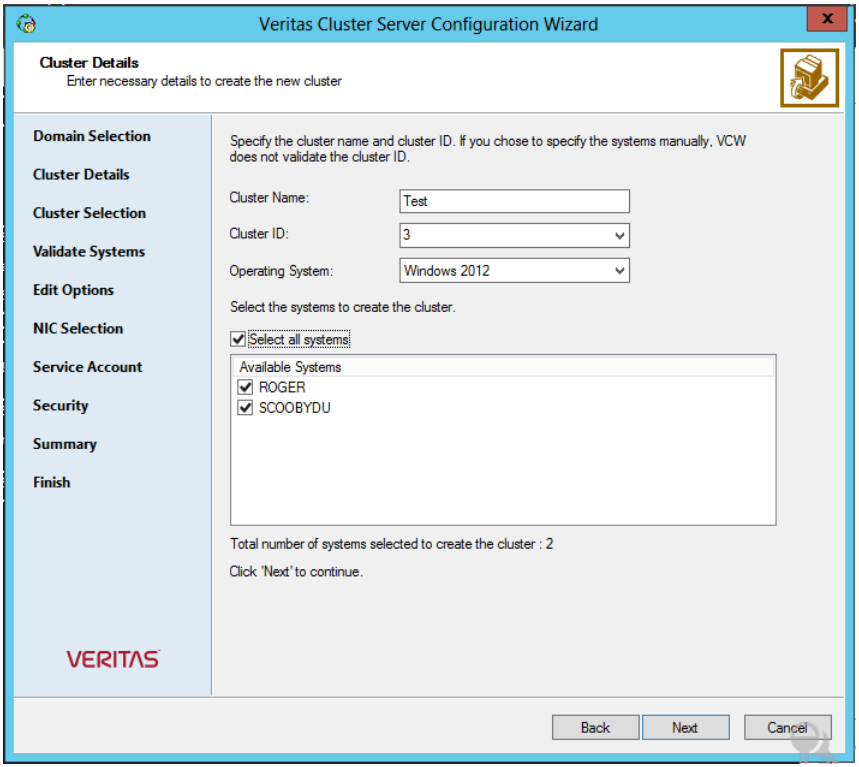

Specify the cluster details as follows:

Cluster Name	Type a name for the new cluster. Veritas 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 65535. **Note:** 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. All the systems in the cluster must have the same operating system and architecture. For example, you cannot configure a Windows Server 2008 R2 system and a Windows Server 2012 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 step 9, proceed to the next step. Otherwise, proceed to step 12.

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 using IPv4 or IPv6 network.

Do one of the following:

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

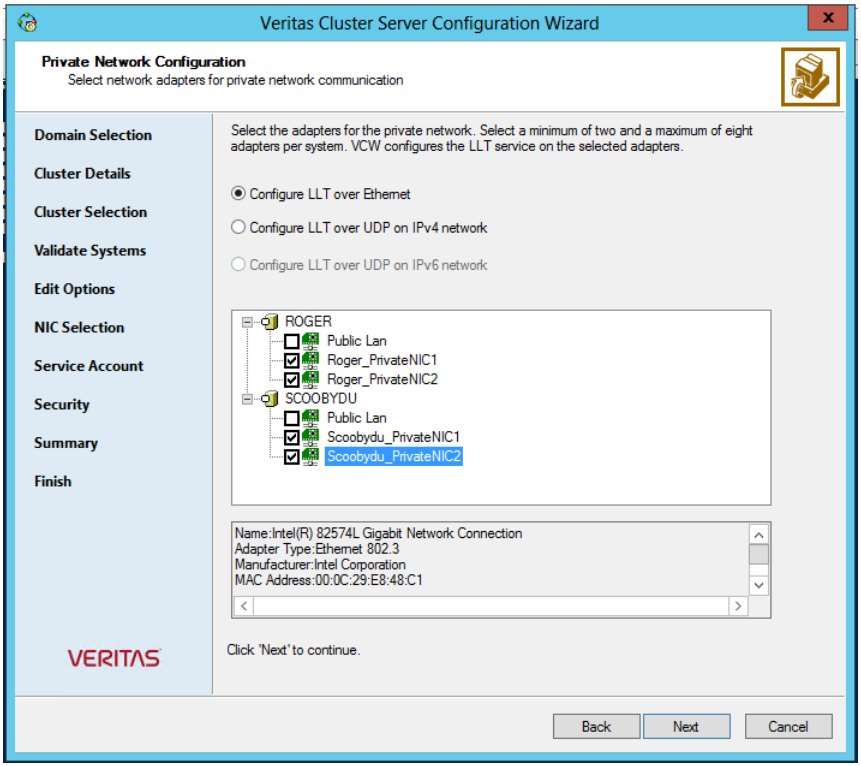

- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.
  Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends 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. Veritas 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 on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.
  The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the VCS private network.
- 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. In case of IPv4, each IP address can be in a different subnet.
  The IP address is used for the VCS private communication over the specified UDP port.
- 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, 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 are used for the private network communication.

12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service.

The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details 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, 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 communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, 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 Cluster Manager (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.

  Veritas recommends that you specify a new user name and password.

- To use the single sign-on feature, click **Use Single Sign-on**.

  In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.

  The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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 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.

16 On the Cluster Service Components panel, select the components to be configured in the ClusterService group and then click **Next**.

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 354.
- 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.

  Note: Configure GCO using this wizard only if you are configuring for disaster recovery but not using the DR wizard.

  Alternatively, consider configuring GCO using the DR wizard.

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 then 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 then click **Next**.

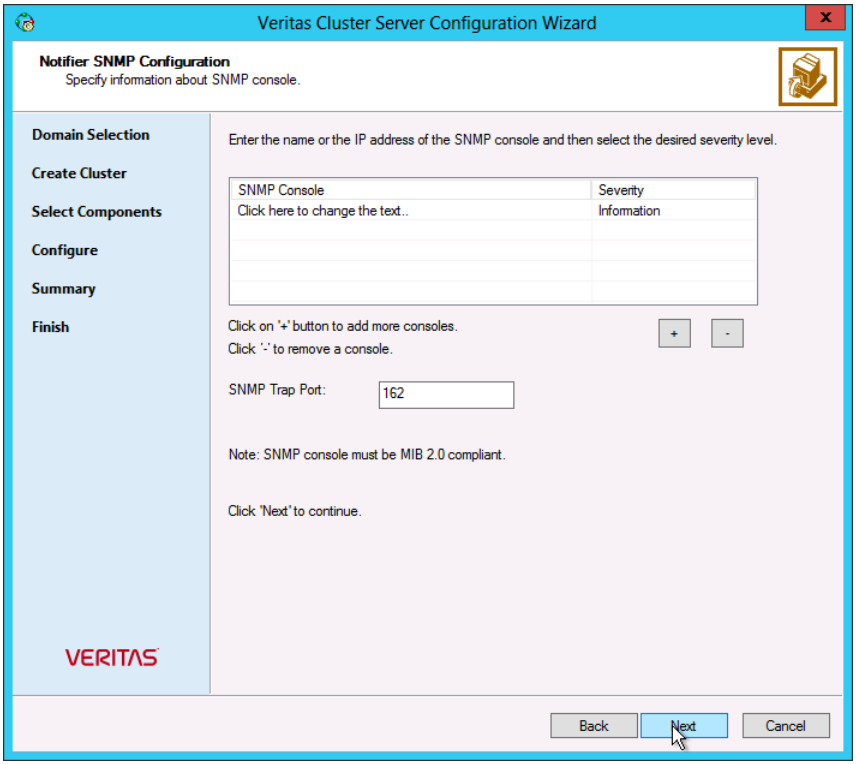

Do the following:

- 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 the **+** icon to add a field; click the **-** icon 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 then click **Next**.

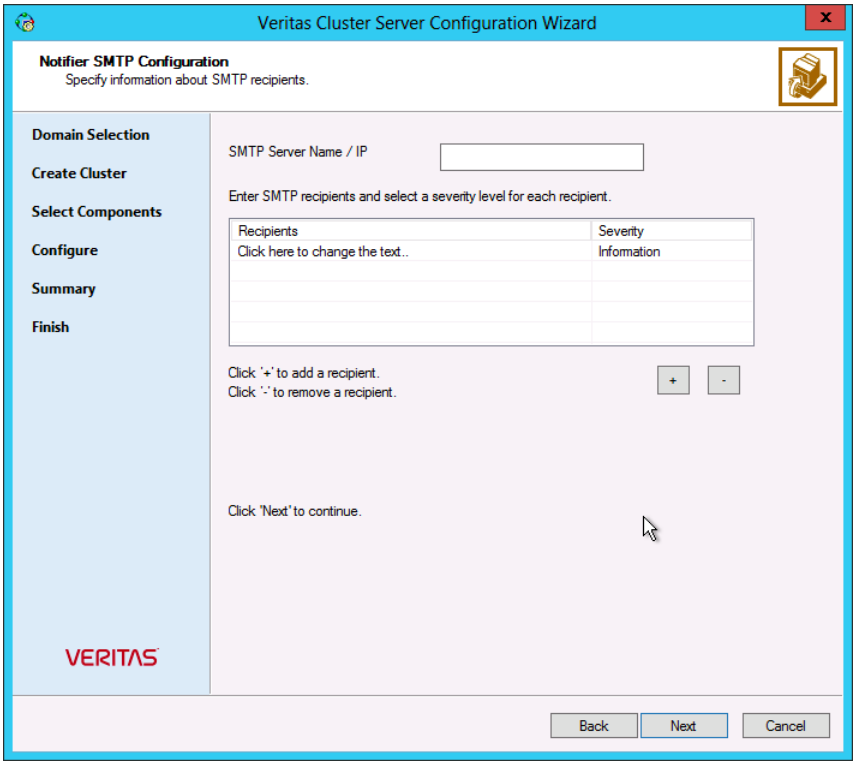

Do the following:

- 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 the **+** icon to add fields; click the **-** icon to remove a field.

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

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.
6. Click **Finish** to exit the wizard.

Creating disk groups and volumes

Use the Veritas Enterprise Administrator (VEA) console to create disk groups and dynamic volumes on the cluster storage, which in a campus cluster consists of two storage arrays.

Configuring cluster disk groups and volumes is covered in the following topics:

- See "About cluster disk groups and volumes" on page 210.
- See "Example disk group and volume configuration in campus cluster" on page 211.
- See "Considerations when creating disks and volumes for campus clusters" on page 506.
- See "Viewing the available disk storage" on page 213.
- See "Creating a dynamic disk group" on page 508.
- See "Adding disks to campus cluster sites" on page 509.
- See "Creating volumes for campus clusters" on page 510.

About cluster disk groups and volumes

SFW uses disk groups to organize disks or LUNs for management purposes. A dynamic disk group is a collection of disks that is imported or deported as a single unit. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

Note: You create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in a high-availability cluster by first deporting the cluster disk group from the current node and then importing it on the desired node. In a campus cluster, the volumes are mirrored across the storage arrays.

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). For more information, see the *Storage Foundation Administrator's Guide*.

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

Example disk group and volume configuration in campus cluster

The illustration that follows shows a VCS campus cluster configuration of disks. For campus clusters, each disk group must contain an equal number of disks on each site. 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 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.

Note: 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

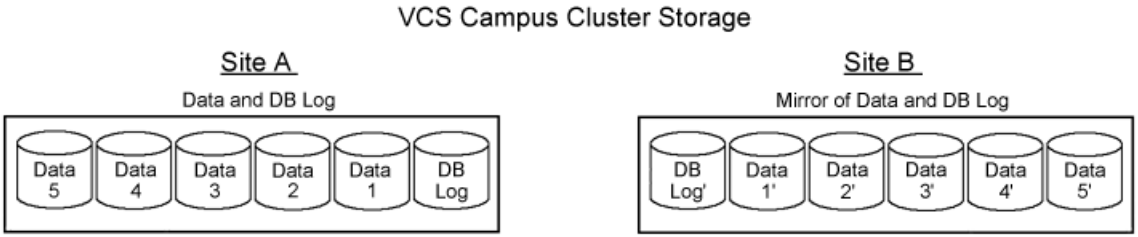

Considerations when creating disks and volumes for campus clusters

When you create the disk groups for a campus cluster, ensure that each disk group has the same number of disks on each physical site. You create each volume as 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.

Veritas recommends using the SFW site-aware allocation feature for campus cluster storage. Site-aware allocation can ensure that site boundary limits are maintained for operations like volume grow, subdisk move, and disk relocation.

Enabling site-aware allocation for campus clusters requires the following steps in the VEA:

- After creating the disk groups, you tag the disks with site names to enable site-aware allocation. This is a separate operation, referred to in the VEA as adding disks to a site.

  As an example, say you had a disk group with four disks. Disk1 and Disk2 are physically located on Site A. Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to "site_a" and add Disk3 and Disk4 to "site_b".

- During volume creation, you specify the volume site type as Site Separated. This ensures that the volume is restricted to the disks on the selected site.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

For more information on site-aware allocation, refer to the *Storage Foundation Administrator's Guide*.

When you create the volumes for a campus cluster, consider the following:

- During disk selection, configure the volume as "Site Separated" and select the two sites of the campus cluster from the site list.
- For volume attributes, select the "mirrored" and "mirrored across enclosures" options.
- Veritas recommends using either simple mirrored (concatenated) or striped mirrored options 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.
- During the volume creation procedure for Site Separated volumes, you can only create as many mirrors as there are sites. However, once volume creation is complete, you can add additional mirrors if desired.
- Choosing "Mirrored" and the "mirrored across" option without having two enclosures that meet requirements causes new volume creation to fail.
- You cannot selecting RAID-5 for mirroring.
- Selecting "stripe across enclosures" is not recommended because then you need four enclosures, instead of two.
- Logging can slow performance.

Viewing the available disk storage

Before creating disk groups and volumes you may want to view available disk storage.

To view the available disk storage

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.

   On Windows 2012 operating systems, use the **Apps** menu.

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**.

4. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

5. 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 disk group

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.

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

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

To create a dynamic disk group

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:
   - Enter the disk group name (for example, **DG1**).
   - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
   - 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.
   - Click **Next**.

   Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).
7. Click **Next** to accept the confirmation screen with the selected disks.
8. Click **Finish** to create the dynamic disk group.

Adding disks to campus cluster sites

For campus cluster storage, Veritas recommends using Storage Foundation (SFW) site-aware allocation. To enable site-aware allocation, you assign a site name to

disks after they are added to a disk group. In the VEA assigning a site name is referred to as adding disks to a site.

For example, Disk1 and Disk2 are physically located on Site A and Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to site_a and add Disk3 and Disk4 to site_b.

To add disks to a site

1. From the VEA console, right-click a disk that needs to be added to a site and select **Add Disk to Site**.

   Disks must be part of a dynamic disk group in order to add them to a site.

2. In the Add Disk to a Site screen, choose one of the following:

   - Choose **Select a new site** and specify a new site name.
     The site name can include any alphanumeric value and valid characters like the period (.), dash (-), and underscore (_). It cannot exceed 31 characters. Site names are case insensitive; all names are converted to lowercase.
   - Choose **Available Sites** and select a site from the list.

3. From the **Available Disks** column, select the disk or disks to add to the specified site.

4. Click **OK**.

Creating volumes for campus clusters

This section will guide you through the process of creating a volume on a dynamic disk group for a campus cluster.

For creating volumes for other types of clusters:

- See "Creating dynamic volumes" on page 475.

Use the following procedure to create dynamic volumes for a campus cluster.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

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

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 as follows:

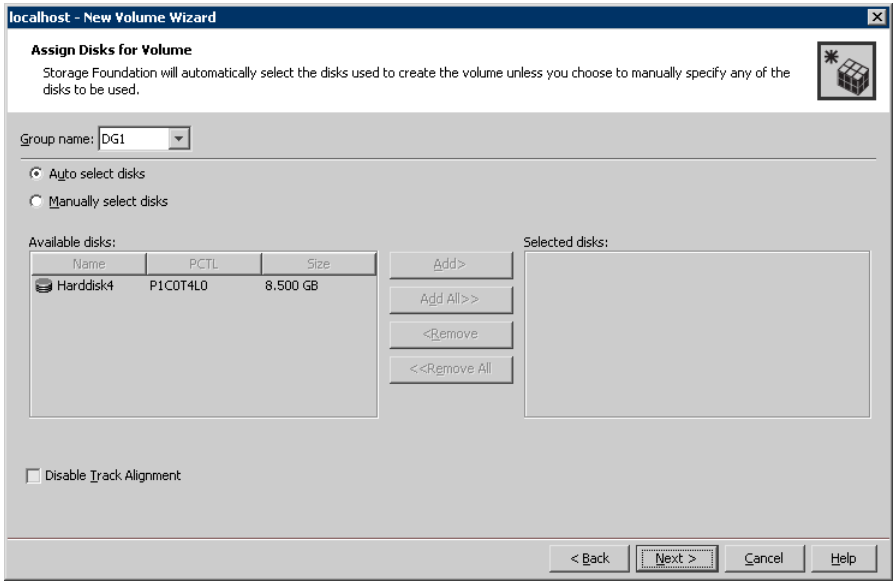

Group name	Make sure the appropriate disk group is selected.
Site preference	Select the **Site Separated** option.
Select site from	Select the campus cluster sites. Press **CTRL** to select multiple sites. **Note:** If no sites are listed, the disks have not yet been added to a site.

Auto select disks	Automatic disk selection 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 picks two disks (one from each array) with the most space.
Manually select disks	If you manually select disks, use the **Add** and **Remove** buttons to move the appropriate disks to the **Selected disks** list.
Disable Track Alignment	You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling track alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

Click **Next**.

7 Specify the volume attributes as follows:

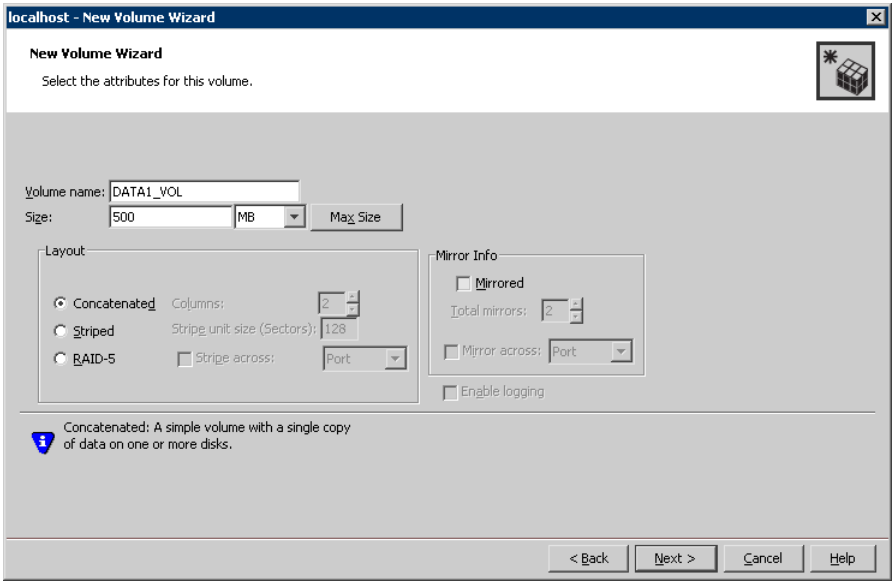

Volume name	Specify a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
Size	Specify a size for the volume. If you click **Max Size**, the **Size** box shows the maximum possible volume size for that layout in the dynamic disk group.
Layout	Ensure that the **Mirrored** checkbox is selected. Select either the **Concatenated** or **Striped** layout type. 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.
Mirror Info	Click **Mirror across** and select **Enclosures** from the drop-down list. When creating a site separated volume, as required for campus clusters, the number of mirrors must correspond to the number of sites. If needed, you can add more mirrors after creating the volume.
Enable logging	Verify that this option is not selected.

Click **Next**.

8 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.

Click **Next**.

9 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**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes as needed.

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.
  See "Deporting and importing a disk group in a campus cluster" on page 223.
- 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.
  See "Deporting and importing a disk group in a campus cluster" on page 223.
- 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. After the application is installed, restart the service.

About configuring a File Share server role in a campus cluster

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.

About configuring a Print Share server role in a campus cluster

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.

About installing and configuring the IIS application in a campus cluster

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.

About installing additional applications in a campus cluster

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 in a campus cluster

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 > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, use the **Apps** menu).

   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 > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, use the **Apps** menu).

   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.
   - 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 consitutes 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 "Configuring the service group" on page 110.

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:

- See "About configuring file shares" on page 110.
- See "About configuring print shares" on page 124.
- See "About configuring IIS sites" on page 132.
- See "Modifying an application service group" on page 170.
- See "About configuring the Oracle service group using the wizard " on page 156.

Verifying the cluster configuration

Simulating a failover is an important part of configuration testing. 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, as follows:

- 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 that you selected to switch to in the first step.
3. To move all the resources back to the original node, repeat the first step of this procedure 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, perform these steps sequentially:
   - Restart the node that you shut down in the first step.
   - 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.

Section 5

Replicated Data Clusters

- Chapter 11. Introduction to Replicated Data Clusters
- Chapter 12. Deploying Replicated Data Clusters: New application installation

Chapter 11

Introduction to Replicated Data Clusters

This chapter includes the following topics:

- About Replicated Data Clusters
- How VCS Replicated Data Clusters work
- Setting up a Replicated Data Cluster configuration
- Migrating the service group

About Replicated Data Clusters

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage (VMDg) or non-shared storage (VMNSDg), 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 Volume Replicator.

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 Volume Replicator 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG 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: Volume Replicator 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 (VMDg) or non-shared storage (VMNSDg) 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.

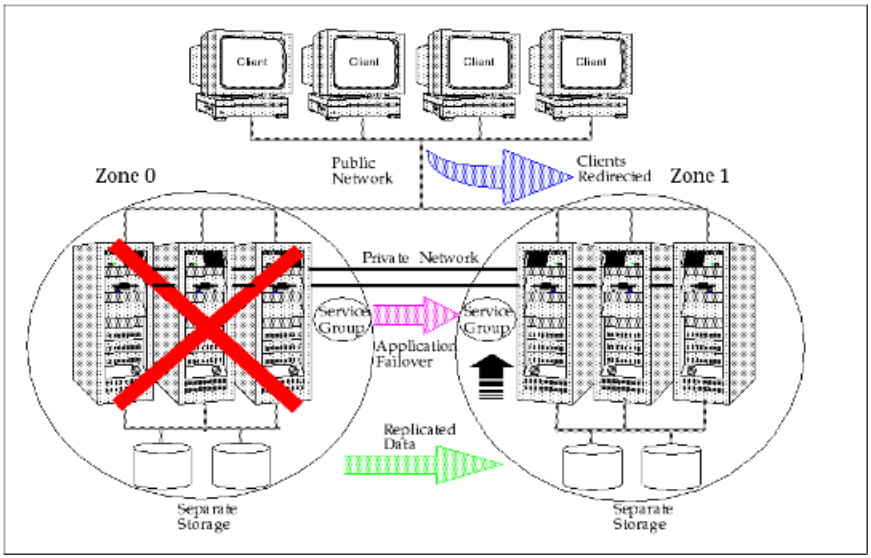

If the system or application fails, VCS attempts to fail over the application service group to another system within the same RDC system zone. However, if VCS cannot find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone (zone 1). VCS also redirects clients once the application is online on the new location.

While this example required using shared storage, you can also set up an RDC cluster that uses non-shared storage. This involves installing and configuring the application on a single system in each of the RDC zones. The application data is located on the local disks attached to the system within each RDC zone. The data is replicated between the systems across the RDC zones to ensure concurrency.

The application service group is online on the single node in the primary RDC zone (Zone 0). In the event of a failure, VCS switches the service group to the node in the secondary RDC zone (Zone 1). Data replication ensures that the application is able to successfully handle client requests from the new node.

The following figure shows failover in a replicated data cluster using non-shared storage.

Figure 11-1 Failover in a replicated data cluster using non-shared storage

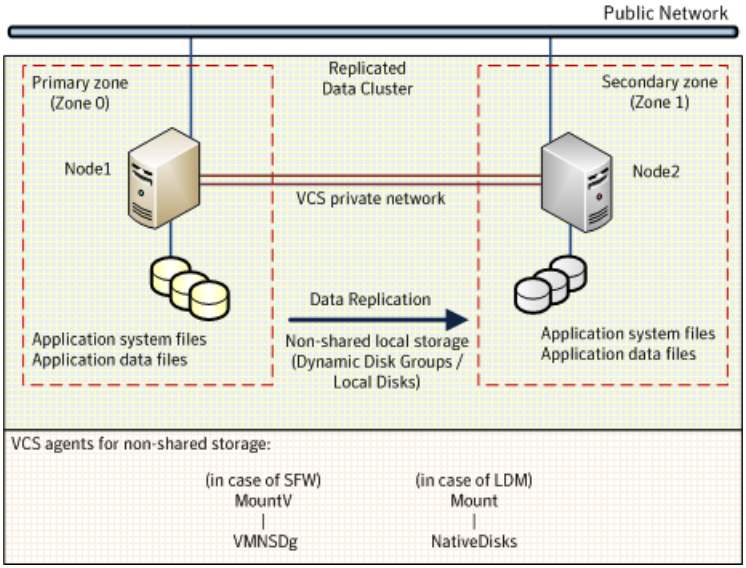

Setting up a Replicated Data Cluster configuration

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:

- See "Setting up replication" on page 231.
- See "Configuring the service groups" on page 232.

Setting up replication

Set up replication between the shared disk groups. Use Volume Replicator to group the shared data volumes into a Replicated Volume Group, and creating the Volume Replicator 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 Volume Replicator 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.

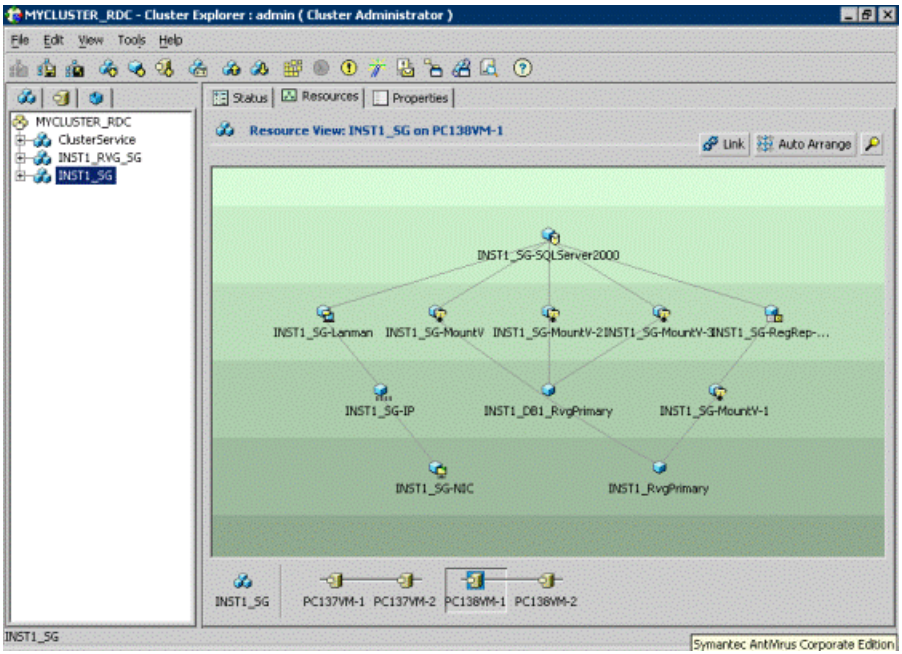

The following screen from the VCS Cluster Manager (Java Console) depicts a typical replication service group (RVG) configuration, again using SQL Server as an example:

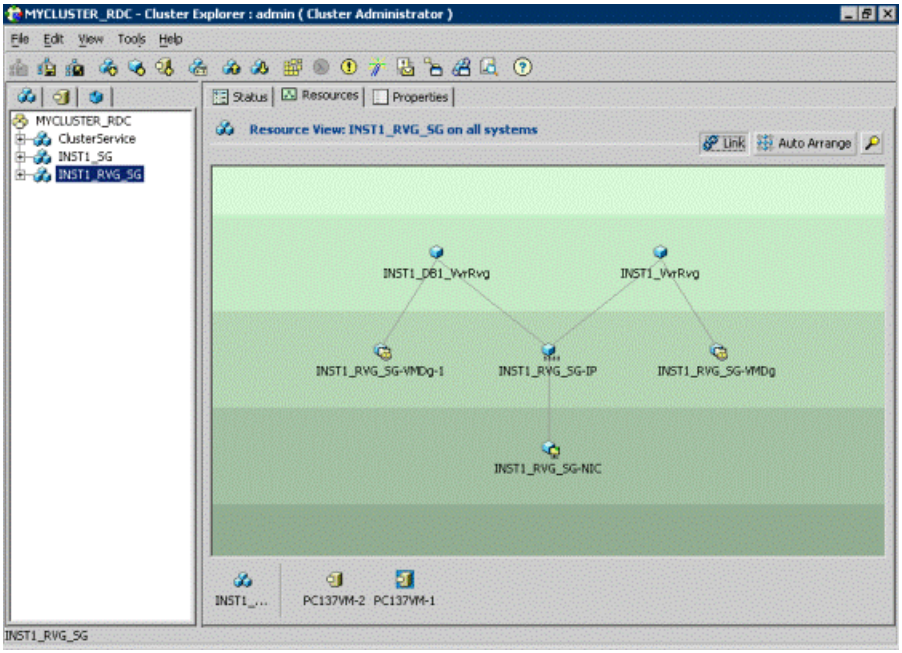

Migrating the service group

In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the shared storage (VMDg) or non-shared storage (VMNSDg). 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 Volume Replicator 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 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**.

Chapter 12

Deploying Replicated Data Clusters: New application installation

This chapter includes the following topics:

- Tasks for a new replicated data cluster installation—additional applications
- Notes and recommendations for cluster and application configuration
- Sample configuration
- Configuring the storage hardware and network
- About installing the Veritas InfoScale products
- Setting up security for Volume Replicator
- Configuring the cluster using the Cluster Configuration Wizard
- Configuring disk groups and volumes
- Installing and configuring the application or server role
- Configuring the service group
- Creating the primary system zone for the application service group
- Verifying the cluster configuration
- Creating a parallel environment in the secondary zone
- Adding nodes to a cluster
- Creating the Replicated Data Sets with the wizard

- Configuring a RVG service group for replication
- Setting a dependency between the service groups
- Adding the nodes from the secondary zone to the RDC
- Verifying the RDC configuration
- Additional instructions for GCO disaster recovery

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 Volume Replicator, see the *Volume Replicator Administrator's Guide*.

The following table outlines the high-level objectives for implementing the configuration and the tasks for each objective.

Table 12-1 Task List: New RDC configuration

Descriptions	Tasks
See "Sample configuration" on page 241.	■ Reviewing the sample configuration
See "Configuring the storage hardware and network" on page 335.	■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which the application will be installed
See "About installing the Veritas InfoScale products" on page 503.	■ Verifying the driver signing option for the system ■ Refer to the *Veritas InfoScale Installation and Upgrade Guide*.
See "Setting up security for Volume Replicator" on page 531.	Using the Volume Replicator Security Service Configuration wizard to configure the VxSAS service for Volume Replicator

Table 12-1 Task List: New RDC configuration *(continued)*

Descriptions	Tasks
See "Configuring the cluster using the Cluster Configuration Wizard" on page 345.	■ 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
See "Configuring disk groups and volumes" on page 258.	■ Planning your storage layout ■ Creating disk groups ■ Creating volumes ■ Managing disk groups and volumes
See "Installing and configuring the application or server role" on page 266.	Installing and configuring the application or server role on the cluster nodes
See "Configuring the service group" on page 268.	■ Using the applicable wizard to create and configure the VCS service group ■ Creating the application service group manually using templates from the Cluster Manager (Java Console) (if using a non-shared storage configuration) ■ Bringing the service group online
See "Creating the primary system zone for the application service group" on page 268.	In the VCS console, selecting the service group and configuring the primary zone nodes as zone 0
See "Verifying the cluster configuration" on page 269.	■ Simulating failover ■ Switching online nodes
See "Creating a parallel environment in the secondary zone" on page 270.	■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage Installing InfoScale Enterprise ■ 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

Table 12-1 Task List: New RDC configuration *(continued)*

Descriptions	Tasks
See "Creating the Replicated Data Sets with the wizard" on page 539.	Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
See "Configuring a RVG service group for replication" on page 288.	■ Creating a Replicated Volume Group (RVG) service group ■ Configuring the RVG service group
See "Setting a dependency between the service groups" on page 299.	Setting up an online local hard dependency of the application service group (the parent) on the RVG service group (the child)
See "Adding the nodes from the secondary zone to the application service group" on page 307.	■ 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
See "Verifying the RDC configuration" on page 312.	Verifying that failover occurs first within zones and then from the primary to the secondary zone

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
  https://sort.veritas.com/documents

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

Refer to the following Microsoft knowledge database article for more details:

http://technet.microsoft.com/en-us/library/dd184075.aspx

- 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 *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. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

  Static IP addresses are required for the following purposes:

  - One static IP address 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 Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
  - For Volume Replicator 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 SFW HA 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 *Cluster Server Bundled Agents Reference Guide*.

- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- 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.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.
- For a Replicated Data Cluster, install only in a single domain.
- 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)
- This is applicable for a Replicated Data Cluster configuration.
  This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, 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.

- To configure a RDC cluster, you need to 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.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.
VCS agents, wizards, and other components	VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above). The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

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.

Zone	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	Disk group names
	FS_REPLOG	Replicator log volume required by Volume Replicator
Secondary zone **Note:** All the other parameters are the same as on the primary zone.	SYSTEM3 & SYSTEM4	First and second nodes of the secondary zone
RDS and Volume Replicator 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.

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, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
   - Veritas 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.

To verify the DNS settings and binding order for all systems

1. Open the Control Panel by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.
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, click **Adapter settings**.
4. Ensure the public network adapter is the first bound adapter by following these steps sequentially:
   - In the Network Connections window, click **Advanced > 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)** or **Internet Protocol Version 6 (TCP/IPv6)** check box, depending on which protocol your network is using.
   - 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 that the **Register this connection's address in DNS** check box is selected.
12. Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.
13. Click **OK**.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Setting up security for Volume Replicator

If you use Volume Replicator for replication, you must configure the Veritas Volume Replicator Security Service (VxSAS) on all the cluster nodes.

In a Replicated Data Cluster environment, you must configure the service on all the nodes in the primary zone as well as the secondary zone.

For details on this required service, see the *Volume Replicator Administrator's Guide*.

After you install InfoScale Storage or InfoScale Enterprise, launch the Veritas Volume Replicator Security Service Configuration Wizard. This wizard lets you complete the Volume Replicator security service configuration.

To do so, launch the wizard after you install InfoScale Enterprise on both the primary and secondary nodes. Then, when you run the wizard, you can specify the primary and secondary sites in one step.

Prerequisites for configuring VxSAS

- The wizard requires you to be logged on with administrative privileges.
- The account that 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.
- The systems on which you want to configure VxSAS must be accessible from the local system.

To configure VxSAS

1. Launch the Veritas Volume replicator Security Service Configuration Wizard from **Start > All Programs > Veritas > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   Optionally, run `vxsascfg.exe` from the command prompt of the required machine.

2. Read the information provided on the Welcome page and click **Next**.

3. Complete the Account Information panel as follows:

Account name (domain\account)	Enter the administrative account name.
Password	Specify a password

   If you have already configured VxSAS for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring VxSAS on the other hosts.

   Click **Next**.

4. On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 lets you specify the domain name. Click **Add** to add the name to the Selected domains list.

   Click **Next**.

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 VxSAS.

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.

When configuring VxSAS in a firewall setup, the VxSAS wizard may not be able to configure the machines that are across the firewall, although the Host Selection dialog may list these nodes. In this case, configure VxSAS locally on the machines that are across the firewall.

Click **Back** to change any information you had provided earlier.

7 Click **Finish** to exit the wizard.

Configuring the cluster using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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 VCW to remove the node from the cluster, rename the system, and then run VCW again to add that system to the cluster.

Note the following prerequisites before you proceed:

- The required network adapters, and SCSI controllers are installed and connected to each system.

  To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.

- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.

- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

  Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.

- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list.

  When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.

- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.

- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
  Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

1. Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
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.

   To discover information about all systems and users in the domain, do the following:

   - Clear **Specify systems and users manually**.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - Select **Specify systems and users manually**.
     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. 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.

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**.

   Select the 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.

   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.
   - Product is either not installed or there is a version mismatch.

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

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

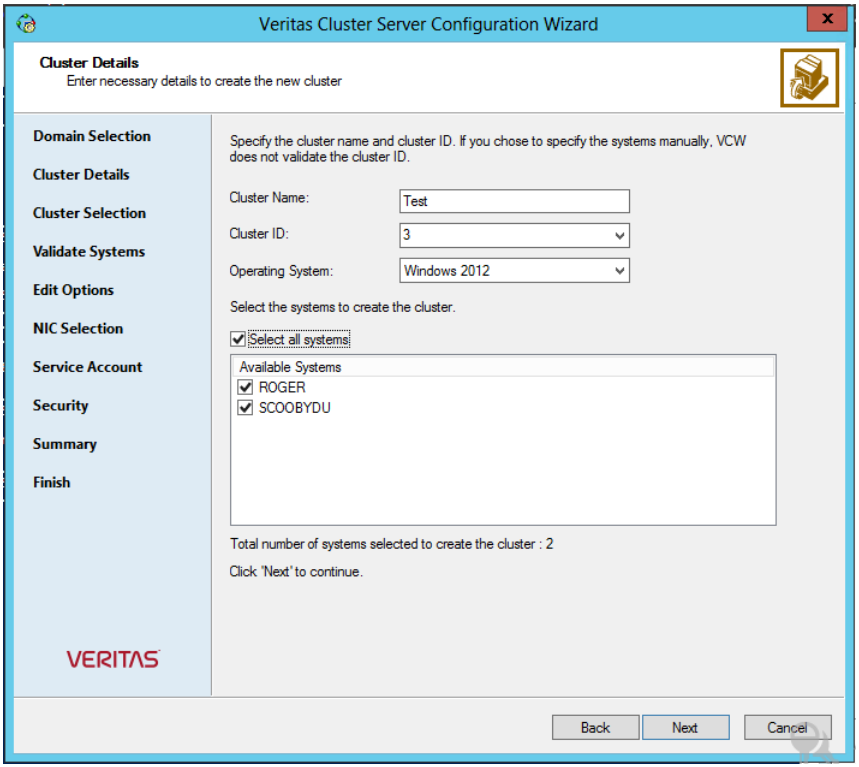

Specify the cluster details as follows:

Cluster Name	Type a name for the new cluster. Veritas 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 65535. **Note:** 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. All the systems in the cluster must have the same operating system and architecture. For example, you cannot configure a Windows Server 2008 R2 system and a Windows Server 2012 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 step 9, proceed to the next step. Otherwise, proceed to step 12.

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 using IPv4 or IPv6 network.

Do one of the following:

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

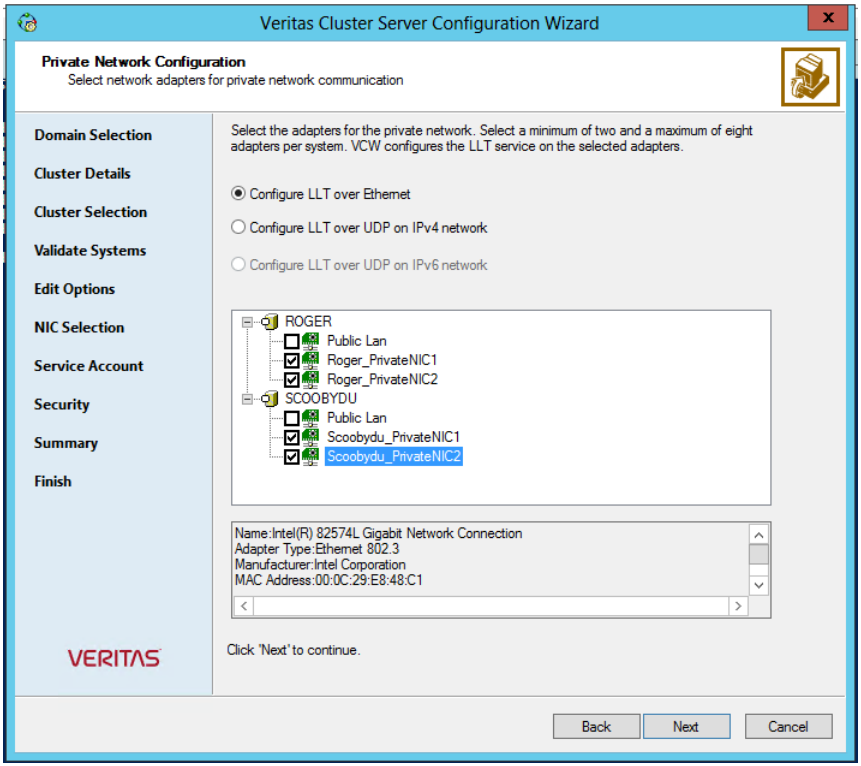

- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.
  Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends 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. Veritas 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 on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.
  The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the VCS private network.
- 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. In case of IPv4, each IP address can be in a different subnet.
  The IP address is used for the VCS private communication over the specified UDP port.
- 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, 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 are used for the private network communication.

12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service.

The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details 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, 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 communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, 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 Cluster Manager (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.

  Veritas recommends that you specify a new user name and password.

- To use the single sign-on feature, click **Use Single Sign-on**.

  In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.

  The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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 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.

16 On the Cluster Service Components panel, select the components to be configured in the ClusterService group and then click **Next**.

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 354.
- 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.

Note: Configure GCO using this wizard only if you are configuring for disaster recovery but not using the DR wizard.

Alternatively, consider configuring GCO using the DR wizard.

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 then 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 then click **Next**.

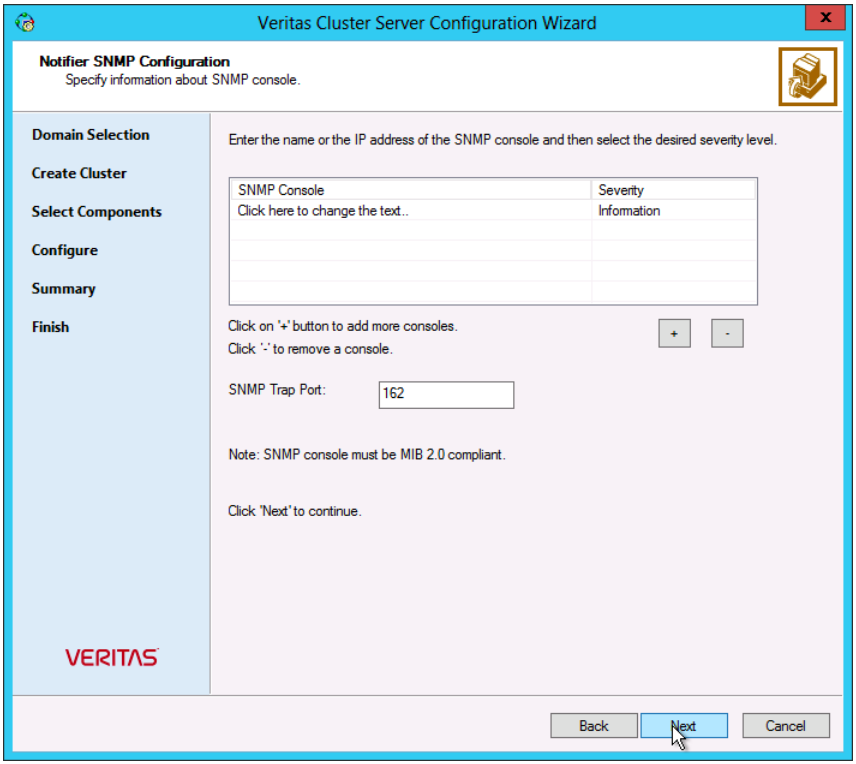

Do the following:

- 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 the **+** icon to add a field; click the **-** icon 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 then click **Next**.

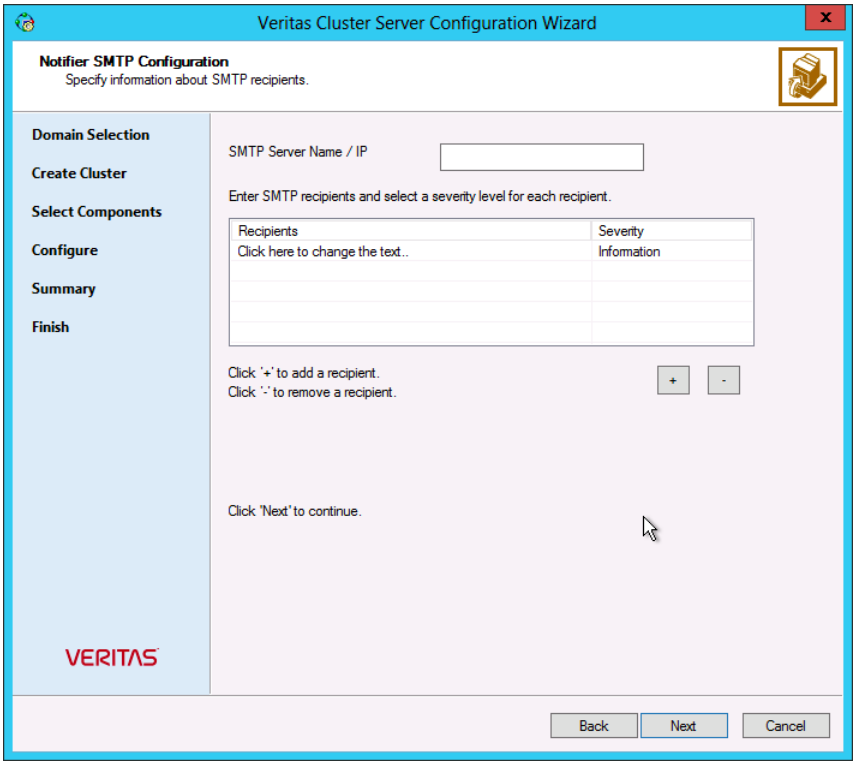

Do the following:

- 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 the **+** icon to add fields; click the **-** icon to remove a field.

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

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.
6. 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 Storage Foundation to create disk groups and dynamic volumes for the application on the shared storage (VMDg) or non-shared storage (VMNSDg).

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**).

For more information, see the *Storage Foundation Administrator's Guide*.

Configuring disk groups and volumes involves the following tasks:

- See "Planning disk groups and volumes" on page 258.
- See "Creating a dynamic disk group" on page 508.
- See "Creating dynamic volumes" on page 475.
- See "About managing disk groups and volumes" on page 337.

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:

- See "Planning your File Share storage" on page 259.

- See "Planning your Print Share storage" on page 259.
- See "Planning your IIS storage" on page 259.
- See "Planning your storage for additional applications" on page 260.

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 (VMDg) or non-shared storage (VMNSDg).
- When you configure a new set up, create the disk groups and volumes on the shared storage (VMDg) or non-shared storage (VMNSDg) first, then create the directory structure for the file shares.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage (VMDg) or non-shared storage (VMNSDg) using the practices recommended by Microsoft.

Planning your Print Share storage

You must create separate volumes on shared storage for the print spooler and for the RegRep volume that contains the list of registry keys that must be replicated between the cluster systems.

Note: Veritas recommends that you place the RegRep and spooler data in separate directories on separate volumes on the shared storage.

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 (VMDg) or non-shared storage (VMNSDg).
- For a new IIS installation, the directory for the web sites must be created on volumes on the shared storage (VMDg) or non-shared storage (VMNSDg).
- For existing web sites, stop the sites and then move the web site content to volumes on the shared storage (VMDg) or non-shared storage (VMNSDg). You must also reconfigure the home directory location for the web site in IIS and then restart the web site again.

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, Veritas 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 system drive of the server. Data files and other related files, such as logs, are placed on the shared storage (VMDg) or non-shared storage (VMNSDg). 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 Volume Replicator, 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 a dynamic disk group

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.

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

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

To create a dynamic disk group

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:
   - Enter the disk group name (for example, **DG1**).
   - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
   - 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.
   - Click **Next**.

Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the dynamic 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.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

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

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. For Site Preference, leave the setting as **Siteless** (the default).
   - Automatic disk selection is the default setting. To manually select the disks, click **Manually select disks** 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.

- Click **Next**.

7 Specify the volume attributes.

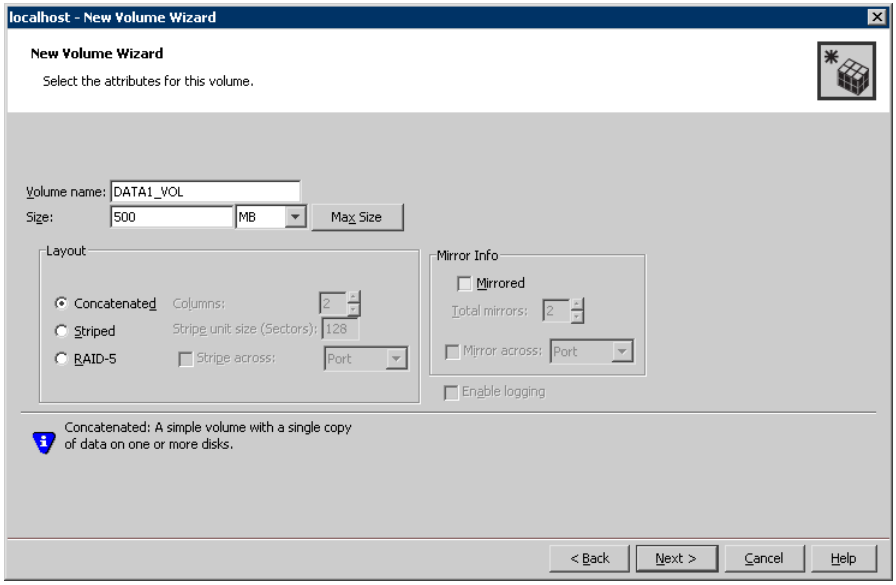

- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns** 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.
- Click **Next**.

8 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 disk.
- Click **Next**.

9 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**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

About managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a cluster dynamic 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.

Managing disk groups and volumes involves the following:

- See “Importing a disk group and mounting a volume” on page 337.
- See “Unmounting a volume and deporting a disk group” on page 338.

Note: (Disaster recovery configurations only) 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 *Storage Foundation Administrator's Guide* for more information.

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:

- See "Configuring a File Share server role" on page 266.
- See "Configuring a Print Share server role" on page 266.
- See "Installing and configuring the IIS application" on page 267.
- See "Installing additional applications" on page 267.

Configuring a File Share server role

When you configure a File Share server role, consider the following:

- Configure the disk group and volumes for the file server shared directory on the shared storage (VMDg) or non-shared storage (VMNSDg).
- When you configure a new setup, create the disk group and volumes on the shared storage (VMDg) or non-shared storage (VMNSDg) first, then create the directory structure for the file shares.
- For an existing configuration that has a file server with shares on the local storage, move these shares to the shared storage (VMDg) or non-shared storage (VMNSDg) using the practices recommended by Microsoft.

Configuring a Print Share server role

When you configure a Print Share server role, consider 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, consider the following:

- Install and configure IIS identically on all nodes hosting the service group. The sites to be monitored must be on the shared storage (VMDg) or non-shared storage (VMNSDg).
- Import the 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 (VMDg) or non-shared storage (VMNSDg) and place the site content in that folder.
- Change the default home directory path for all the IIS sites to be monitored to a location on the shared storage (VMDg) or non-shared storage (VMNSDg). 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 (VMDg) or non-shared storage (VMNSDg). 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.

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 (VMDg) or non-shared storage (VMNSDg).

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:

- See "About configuring file shares" on page 110.
- See "About configuring print shares" on page 124.
- See "About configuring IIS sites" on page 132.
- See "Modifying an application service group" on page 170.

Creating the primary system zone for the application service group

In the application 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 **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.

In case of a non-shared storage configuration, add only the single node to the primary zone.

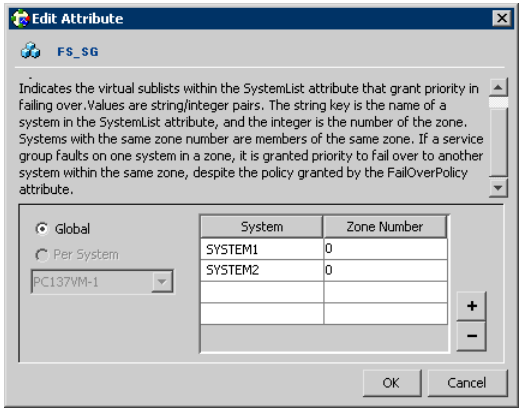

7 Click **OK**.

After setting up the primary system zone, you can verify the service group failover on systems within the primary zone.

See "Verifying the RDC configuration" on page 312.

Verifying the cluster configuration

Simulating a failover is an important part of configuration testing. 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, as follows:

- 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 that you selected to switch to in the first step.
3. To move all the resources back to the original node, repeat the first step of this procedure 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, perform these steps sequentially:
   - Restart the node that you shut down in the first step.
   - 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, in the application service group in the primary zone, take the following resources offline:

- Application resource
- Application virtual IP resource

The remaining resources should be online, including the storage 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 sections to complete the same tasks in the secondary zone (zone 1).

- See "Notes and recommendations for cluster and application configuration" on page 330.
- See "Configuring the storage hardware and network" on page 335.
- See "About installing the Veritas InfoScale products" on page 503.
- See "Setting up security for Volume Replicator" on page 531.
- See "Adding nodes to a cluster" on page 271.
- See "Configuring disk groups and volumes" on page 258.
  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:
  - Disk group name
  - Volume sizes
  - Volume names
  - Drive letters
- See "Installing and configuring the application or server role" on page 266.

Adding nodes to a cluster

If you are setting up a Replicated Data Cluster, use the VCS Cluster Configuration Wizard (VCW) to add the systems in the secondary zone (zone1) to the existing cluster.

You use the VCS Cluster Configuration Wizard (VCW) to add one or more nodes to an existing cluster.

Prerequisites for adding a node to an existing cluster are as follows:

- Verify that the logged-on user has VCS cluster administrator privileges.
- The logged-on user must be a local administrator on the system where you run the wizard.
- Verify that Command Server is running on all nodes in the cluster. Select Services on the Administrative Tools menu and verify that the Veritas Command Server shows that it is started.
- Verify that the high availability daemon (HAD) is running on the node on which you run the wizard. Open the Services window, and verify that the **Veritas High availability engine** service is running.

To add a node to a VCS cluster

1. Start the VCS Cluster Configuration wizard.

   Click **Start > All Programs > Veritas > 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.

   To discover information about all the systems and users in the domain, do the following:

   - Clear the **Specify systems and users manually** check box.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - 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. Otherwise proceed to the next step.

5. On the System Selection panel, complete the following and click **Next**:

   - Type the name of an existing 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.

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.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's 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.

Click **Next** to proceed.

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 (a cluster that is not a secure cluster).

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 use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over Ethernet, do the following:
  - Select the check boxes next to the two NICs to be assigned to the private network.
    Veritas 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 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. Veritas 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, do the following:
  - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Veritas recommends reserving at least two NICs exclusively for the VCS private network. 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.

- 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. In case of IPv4, 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 are 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 credentials 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**.

If you are setting up a Replicated Data Cluster, return to the task list:

See “Creating a parallel environment in the secondary zone” on page 270.

Creating the Replicated Data Sets with the wizard

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.

Configuring Volume Replicator 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.

If you are using IPv6, make sure that you have set the appropriate IP preference for Volume Replicator before configuring Volume Replicator replication. The default setting is IPv4. When you specify host names while configuring replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP protocol Volume Replicator uses to resolve the host names. Use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to set the IP preference.

- Verify that the data volumes are not of the following types, as Volume Replicator does not support these types of volumes:
  - Storage Foundation (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 disk group is imported and the volumes are mounted in the primary and secondary zone.
- Verify that you have set the appropriate IP preference, whether Volume Replicator should use IPv4 or IPv6 addresses. The default setting is IPv4.

  When you specify host names while configuring replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP protocol Volume Replicator uses to resolve the host names. Use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to set the IP preference.
- Configure the VxSAS service if you have not already done so.

  See "Setting up security for Volume Replicator" on page 531.

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.

   Start VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**.

   On Windows 2012 operating systems, from the **Apps** menu in the Start screen.

   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**.

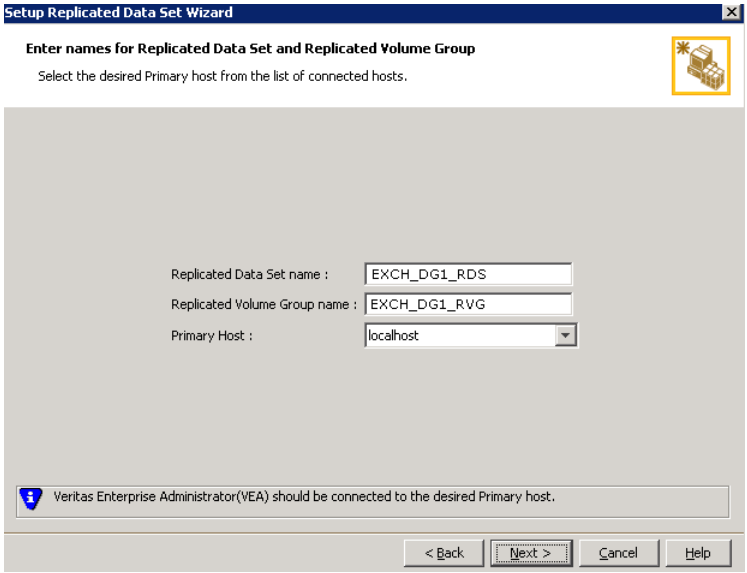

By default, the local host is selected as the Primary Host. To specify a different host name, make sure the required host is connected to the VEA console and select it in the Primary Host list.

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**.

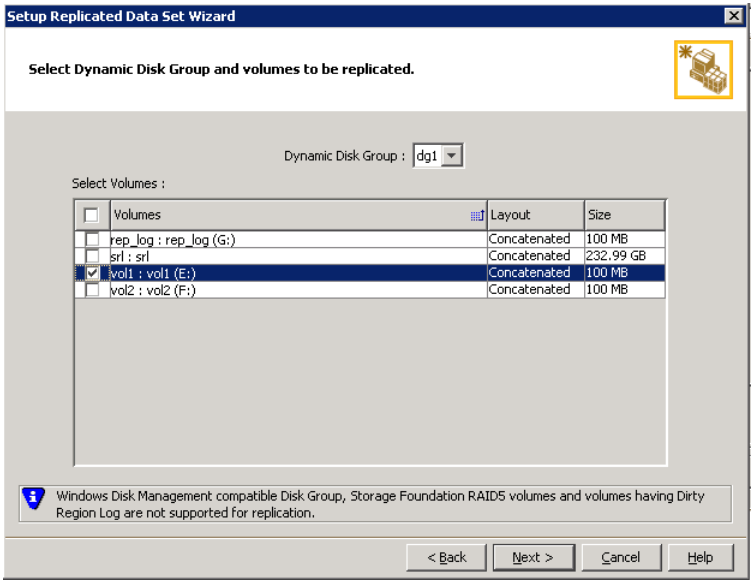

To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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:

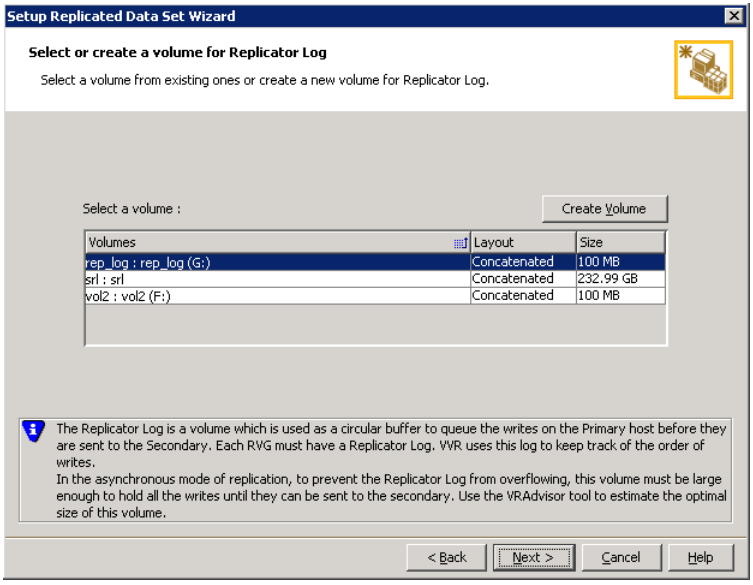

To select an existing volume

- 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 appears:

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 *Storage Foundation Administrator's Guide*.
- Choose the **Select disks automatically** option if you want Volume Replicator 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, Volume Replicator 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 Volume Replicator displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

- The same or larger amount of space as that on the Primary
- Enough space to create volumes with the same layout as on the Primary Otherwise, the RDS setup wizard enables you to create the required volumes manually.
- Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
- Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.

11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

- If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

  Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

  When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

12 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

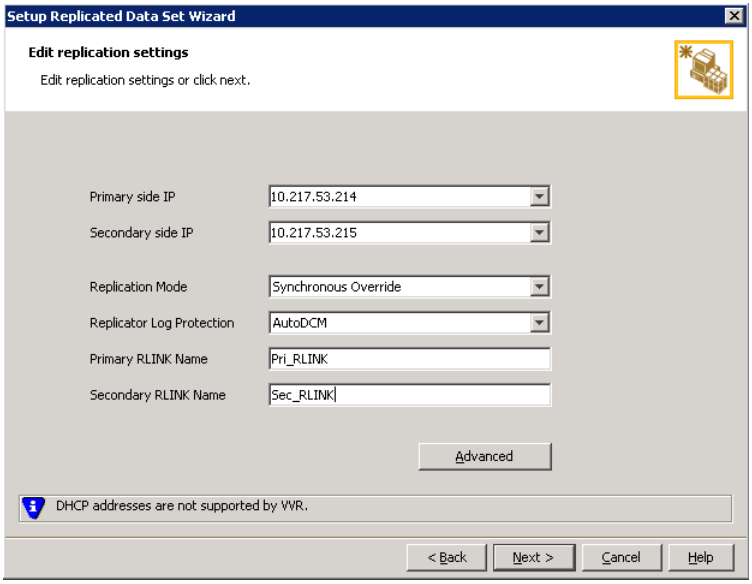

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side	IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode

Select the required mode of replication:

- **Synchronous Override** (default) 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 Volume Replicator updates in the Replicator Log. From there, Volume Replicator 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	■ **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 Volume Replicator 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 Volume Replicator assigns a default name.

- If you want to specify advanced replication settings, click **Advanced**. Edit the replication settings for a secondary host as needed.

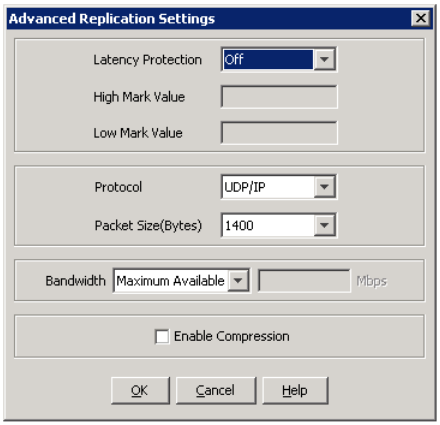

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

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, Volume Replicator 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.
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, Volume Replicator 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**.

13 On the Start Replication page, choose the appropriate option as follows:

- To add the Secondary and start replication immediately, select **Start Replication** with one of the following options:

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, Volume Replicator 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 and ReFS 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 *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.

14 Review the information.

Click **Back** to change any information you had specified.

Otherwise, 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 this procedure for each additional disk group. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name.

Configuring a 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

For more information about service group types, see the *Cluster Server Administrator's Guide*.

Configure the RVG service group resources manually by copying and modifying components of the application service group. Then create new RVG resources and bring them online.

The following table lists the resources in the RVG service group for RDC.

Table 12-2 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg (shared storage) or VMNSDg (non-shared storage) for the disk group	Colume 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.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

To create a 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:

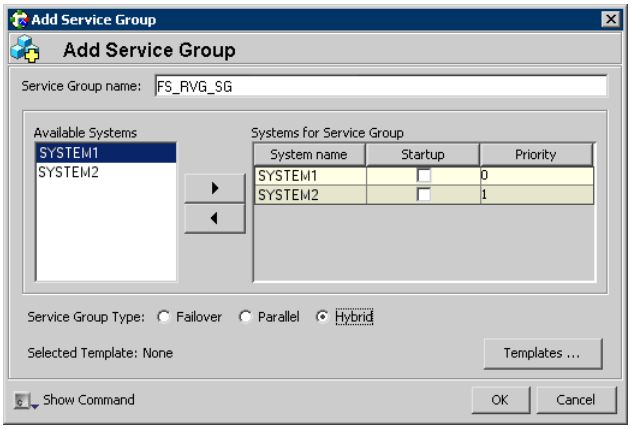

- Enter a name for the service group. Choose a service group name that is meaningful for the type of application you are using. This example adds a FileShare RVG service group called FS_RVG_SG.
- Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.

  Note: If you are creating the RVG service group for a DR configuration in a non-shared storage environment, select **Failover**.

- Click **OK**.

Note: If you are setting up replication in a non-shared storage environment, you can use the replication service group template, **VvrRvgVMNSRVGGroup**, available in the Java Console. For an RDC configuration, ensure that you select the service group type as Hybrid while creating the service group using the Configuration Wizard from Java Console.

Configuring the resources in the RVG service group for RDC replication

Configure the RVG service group's resources manually for RVG by completing the tasks in the following table:

Table 12-3 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 393	See “Configuring the IP and NIC resources” on page 290.
Copy the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resources for the disk groups in the application service group, paste and modify them for the RVG service group. “Configuring the VMDg or VMNSDg resources for the disk groups” on page 394	See “Configuring the VMDg or VMNSDg resources for the disk groups” on page 292.
Create the Volume Replicator RVG resources for the disk group and enter the attributes for the disk group and the replication IP address. “Adding the Volume Replicator RVG resources for the disk groups” on page 397	See “Adding the Volume Replicator RVG resources for the disk groups” on page 294.
Link the Volume Replicator RVG resources to establish the dependencies between the VMDg or VMNSDg resources, the IP resource for replication, and the Volume Replicator RVG resources for the disk group. Configure the RVG service group’s VMDg or VMNSDg resources to point to the disk group that contain the RVGs. “Linking the Volume Replicator RVG resources to establish dependencies” on page 398	See “Linking the Volume Replicator RVG resources to establish dependencies” on page 295.
Delete the VMDg or VMNSDg resources from the application service group, because they depend on the replication and were configured in the RVG service group. “Deleting the VMDg or VMNSDg resource from the application service group” on page 398	See “Deleting the VMDg or VMNSDg resource from the application service group” on page 296.

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC:

Table 12-4 IP and NIC resources

Resource	Attributes to modify
IP	Address
NIC	(none)

Note: In a non-shared storage environment, if you use the Java Console template, **VvrRvgVMNSRVGGroup**, to create the RVG service group, then do not recreate these resources; you modify the attributes of the existing IP and NIC resources in the service group.

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**.
5. In the Name Clashes window, change the names of the IP and NIC resources for the RVG service group and click **OK**.

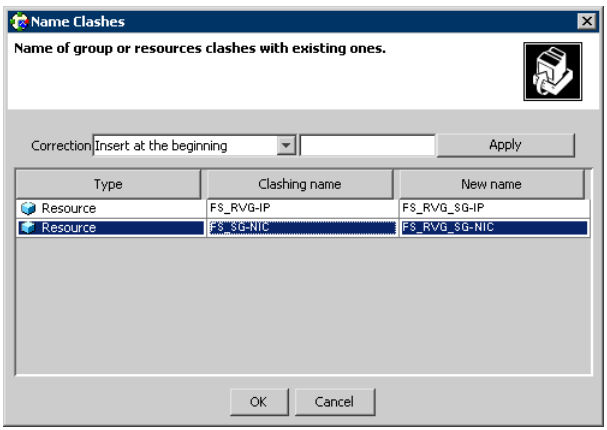

To modify the IP resource and NIC

1. In the Resources tab display area, right-click the IP resource and select **View > Properties View**.
2. In the Properties View window, for the Address attribute, click **Edit**.
3. In the Edit Attribute window, enter the Volume Replicator IP address for the Primary Zone as the scalar value.

   This is the IP address you specified as the Primary side IP address while configuring the Replicated Data Set (RDS) earlier using the RDS wizard.
4. Close the Properties View window.

To enable the IP resource and NIC

1. In the Resources tab display area, right-click the IP resource and select Enabled.
2. In the Resources tab display area, right-click the NIC resource and select Enabled.

Configuring the VMDg or VMNSDg resources for the disk groups

Configuration involves the following tasks:

- Create the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resources in the RVG service group by copying them from the application service group.
- If you are creating a DR configuration in a non-shared storage environment, modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately to ensure the desired failover behavior.

Configure the following attributes in the application service group for the MountV resource.

Resources for the disk group for the application system files:

- VMDg Resource Name
- Volume Name

To create a VMDg or VMNSDg 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 or VMNSDg 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 or VMNSDg resource for the RVG service group.

6 Click **OK**.

Modify the DGGuid attribute for the new disk group resource.

See "Modifying the DGGuid attribute for the new disk group resource in the RVG service group" on page 293.

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 (in case of shared storage) or VMNSDg (in case of non-shared storage) resource that was just created in the RVG service group.

6 Close the Properties View window.

7 Repeat this procedure for any additional MountV resources for the application.

To enable the VMDg or VMNSDg resource in the RVG service group

1 In the left pane, select the RVG service group.

2 In the Resources tab display area, right-click the VMDg or VMNSDg resource and select **Enabled**.

Modifying the DGGuid attribute for the new disk group resource in the RVG service group

To modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service group

1 From the VCS Cluster Explorer, in the left pane, select the RVG service group.

2 In the right pane, select the Resources tab.

3. In the Resources tab display area, right-click the new VMDg or VMNSDg resource and click **View > Properties View**.
4. In the Properties View window, locate the DGGuid attribute and click the edit icon.
5. In the Edit Attribute window, edit the attribute by performing the following actions sequentially:
   - Select **Per System**.
   - From the dropdown list select the first node in the primary zone (Zone 0).
   - In the **Scalar Value** field specify the GUID of the disk group that is imported on the node.
     Run the `VMGetDrive` utility at the command prompt to retrieve the GUID.
   - Repeat the previous two steps, and select a different node from the dropdown list each time. You must specify the GUID separately for each node displayed in the dropdown list.
     In case of a shared storage environment (VMDg resource), if there are multiple nodes in the primary zone, then the disk group GUID will be the same for all systems within the zone. However, the GUID will always be different across zones.
6. In the Properties View window, verify that all nodes in the RDC primary zone have DGGuid values specified.

   Note: If you are creating a DR configuration manually for a non-shared storage environment, you have to modify the DGGuid attribute of the VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately.

7. Close the Properties View window.

Adding the Volume Replicator 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.

Resources for the disk group for the application files:

- VMDgResName
- IPResName

To create the Volume Replicator 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 Volume Replicator 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.

   The RVG name is the name you specified when you created the Replicated Data Set (RDS) earlier using the RDS wizard. You can retrieve the RVG name by running the command `vxprint -VPl`.
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 Volume Replicator 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.

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 or VMNSDg for the application, for example: FS_RVG_SG-VMDg or FS_RVG_SG-VMNSDg

To link the Volume Replicator RVG resources

1. In the left pane, select the RVG service group.
2. Click the Link button in the right pane.
3. To link the VvrRvg resource to the IP resource, click the parent resource, for example FS_RVG_SG_VvrRvg, and then click the child resource, for example FS_RVG_SG_-IP.
4. When prompted to confirm, click **OK**.
5. To link the VvrRvg resource to the VMDg or VMNSDg resource, click the parent resource, for example INST1_DB1_VvrRvg, and then click the child resource, for example INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.
6. When prompted to confirm, click **OK**.
7. 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 or VMNSDg resource from the application service group

The VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) 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 or VMNSDg 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 or VMNSDg 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 or VMNSDg 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. This is the name of the VvrRvg resource in the RVG replication service group.

Configure the following attributes in the application service group for the RVG Primary resources.

Resource for the disk group for the application:

Resource: RVGPrimary

Attribute: 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:

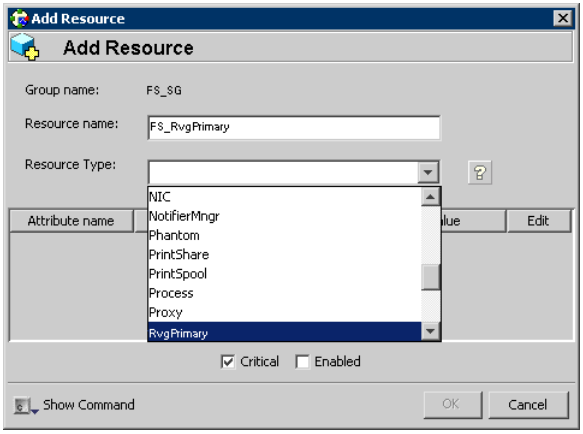

- Enter the Resource Name for the RVG Primary resource for the application disk group.
- Select the Resource Type of RVGPrimary.

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 Volume Replicator RVG resource, and click **OK**.

   This is the name of the VvrRvg resource in the RVG replication service group.

5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.

   See the *Volume Replicator 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, the following table lists the Parent and Child relationship for the FileShare application:

Table 12-5 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 service group

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 **Show All Attributes** button.
4. In the Attributes View, scroll down and select the **SystemZones** attribute, and click the adjacent **Edit** icon.
5. In the Edit Attribute dialog box, click the plus sign (**+**) and enter the systems and the zone number (zone 0) for the primary zone.

   In case of a non-shared storage configuration, add only the single node to the primary zone.
6. 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 online local relationship.
- Select the hard dependency type.
- 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.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the RVG service group.

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 > Veritas > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.

On Windows 2012 operating systems, use the **Apps** menu.

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.
- To enable the service group to automatically come online on one of the systems, select the Include selected systems in the service group's AutoStartList attribute checkbox.
  For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.
- 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 When prompted, click **Yes** to modify the service group.

14 Click **Finish**.

Note: Use the following procedure if the RVG service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the RVG service group using Java Console

1. From VCS Cluster Explorer, in the left pane, right-click the RVG service group and select **View > Properties View**.
2. In the Attributes window, click **Show all attributes**.
3. From the attributes list, select the attribute **SystemList** and click the **Edit** icon.
4. In the Edit Attribute window, edit the SystemList attribute as follows:
   - Click the **+** button to add an empty row.
   - In the **System** field type the cluster node name from the secondary zone.
   - In the **Priority** field type 1.
   - Click **OK**.
5. Close the Attributes windows.

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 **Show All Attributes** button.
4. In the Attributes View, scroll down and select the **SystemZones** attribute, and click the adjacent **Edit** icon.
5. If a message appears indicating that the configuration be changed to read/write, click **Yes**.
6. In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.

   In case of a non-shared storage configuration, add only the single node to the secondary zone.
7. Click **OK** and close the Attributes View window.

Configuring the RVG service group NIC resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource.

Modify the MACAddress attribute of the NIC resource in the RVG service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the RVG service group

1. From the VCS Cluster Explorer, in the left pane, select the RVG service group.
2. In the right pane, select the Resources tab.
3. In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.
4. In the Properties View window, locate the MACAddress attribute and click the edit icon.
5. In the Edit Attribute window, edit the attribute by performing the following actions sequentially:

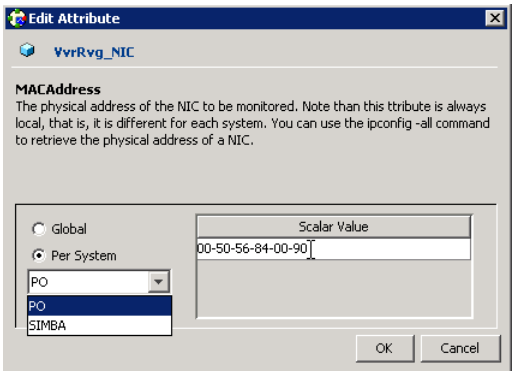

   - Select **Per System**.
   - From the dropdown list, select the node in the RDC primary zone.
   - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
   - From the dropdown list, select the node in the RDC secondary zone.
   - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
   - Click **OK**.
6. In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
7. Close the Properties View window.

Configuring the RVG service group IP resource 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.

Use the following procedure to modify the IP resources.

Note: For IPv6 networks, modify the IPv6 resource.

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. The IP address at the primary zone and the secondary zone should be different.

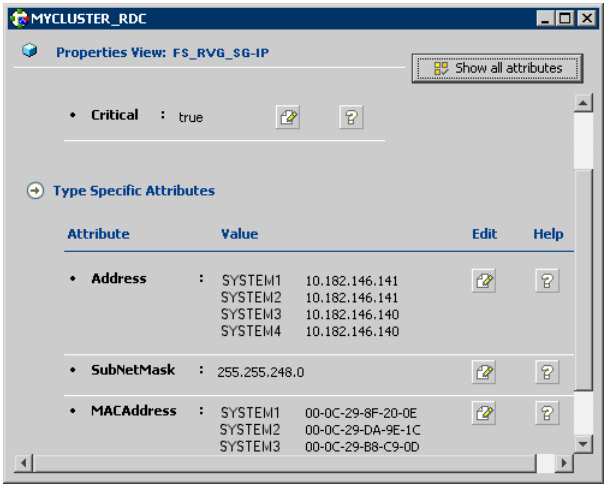

6 This step is applicable only if you are using a non-shared storage environment (VMNSDg agent).

In the Edit Attributes window, edit the MACAddress attribute as follows:

- Select **Per System**.
- From the dropdown list, select the node in the RDC primary zone.
- In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
- From the dropdown list, select the node in the RDC secondary zone.
- In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
- Click **OK**.

7 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.

8 Close the Properties View window.

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.

Configuring the RVG service group VMNSDg resources for fail over

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource. Modify the DGGuid attribute of the VMNSDg resources in the RVG service group to ensure the desired failover behavior in the RDC.

To modify the VMNSDg 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 VMNSDg resource and select **View > Properties View**.

4 In the Edit Attributes window, edit the DGGuid attribute by performing the following actions sequentially:

- Select **Per System**.
- Select the node in the RDC secondary zone.
- In the Scalar Value field, enter the GUID of the dynamic disk group that is imported on the single node in the RDC secondary zone.
- You can retrieve the disk group details running the VMGetDrive utility from the command prompt.
- Click **OK**.

5 In the Properties View window, verify that the DGGuid for the nodes in the primary and secondary zone are different.

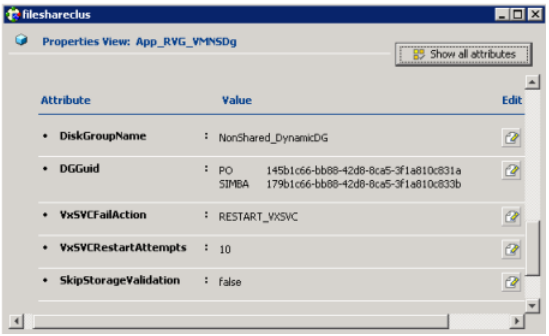

6 Close the Properties View window.

As 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.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the application service group.

Use the following procedure if the service group contains a VMDg resource (shared storage environment).

To add nodes from the secondary zone to the application service group

1 Select **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > File Share Configuration Wizard**.

On Windows 2012 operating systems, use the **Apps** menu.

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. Then, select the File Share service group below this option and click **Next**.
4. In the Service Group Configuration window, review and make the following updates if required:
   - To add nodes from the secondary zone to the application service group, select them in Available Cluster Systems and use the right arrow button to move them to Systems in Priority Order.
   - To change the priority of a system, select the system in the Systems in Priority Order list and click the up and down arrow buttons.
     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.
   - To enable the service group to automatically come online on one of the systems, select the Include selected systems in the service group's AutoStartList attribute checkbox.
     For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.
   - Click **Next**.
5. In the FileShare Configuration window, enter a virtual computer name and select the directories to be shared. Click **Next**.
6. In the Share Permissions window, select the fileshare users and assign them permissions. Click **Next**.
7. 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**.
8. In the Service Group Summary window, review the service group configuration.
   - Press F2 to edit the resource names.
   - To enable all the VMDg resources in the service group for fast failover, select the Enable FastFailOver attribute for all the VMDg resources in the service group checkbox.
     For information about the FastFailOver attribute, see the *Cluster Server Administrator’s Guide*.
   - Click **Next**.

9 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.

10 Click **Finish**.

Use the following procedure if the service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the application service group using Java Console

1 From VCS Cluster Explorer, in the left pane, right-click the application service group and select **View > Properties View**.

2 In the Attributes window, click **Show all attributes**.

3 From the attributes list, select the **SystemList** attribute and click the edit icon.

4 In the Edit Attribute window, edit the SystemList attribute as follows:

- Click the **+** button to add an empty row.
- In the **System** field type the cluster node name from the secondary zone.
- In the **Priority** field type 1.
- Click **OK**.

5 Close the Attributes windows.

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 **Show All Attributes** button.

4 In the Attributes View, scroll down and select the **SystemZones** attribute, and click the adjacent edit icon.

5 If a message appears indicating that the configuration be changed to read/write, click **Yes**.

6. In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.

   In case of a non-shared storage configuration, add only the single node to the secondary zone.

7. Click **OK**.

8. Close the Attributes View window.

Configuring the application service group IP resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the IP resource in the application service group to ensure the desired failover behavior in the RDC.

Note: For IPv6 networks, modify the IPv6 resources.

To modify the IP resource 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 Resources tab.

3. Right-click the IP resource and select **View > Properties View**.

4. In the Edit Attributes window, edit the MACAddress attribute by performing these actions sequentially:

   - Select **Per System**.
   - From the dropdown list, select the node in the RDC primary zone.
   - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
   - From the dropdown list, select the node in the RDC secondary zone.
   - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
   - Click **OK**.

5 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.

6 Close the Properties View window.

Configuring the application service group NIC resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the MACAddress attribute of the NIC resource in the application service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the application service group

1 From the VCS Cluster Explorer, in the left pane, select the application service group.

2 In the right pane, select the Resources tab.

3 In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.

4 In the Properties View window, locate the MACAddress attribute and click the edit icon.

5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:

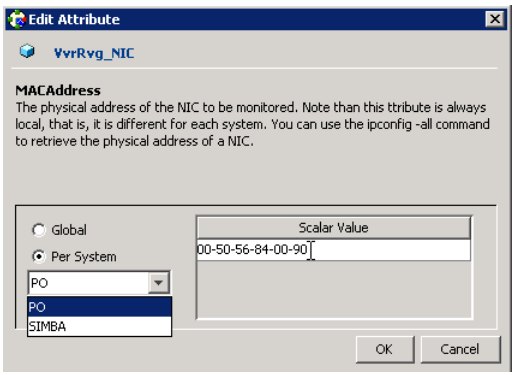

- Select **Per System**.
- From the dropdown list, select the node in the RDC primary zone.
- In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system.

Use the `ipconfig -all` command to retrieve the physical address.

- From the dropdown list, select the node in the RDC secondary zone.
- In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
- Click **OK**.

6. In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
7. Close the Properties 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:

- See "Bringing the service group online" on page 312.
- See "Switching online nodes" on page 312.

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 service group online in the primary zone.

To bring the service group online

1. From VCS Cluster Explorer, in the left pane, right-click the service group (for example, 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 > Veritas > Veritas Cluster Server > Veritas Cluster Manager - Java Console)**.

   On Windows 2012 operating systems, use the **Apps** menu.

2. Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, click **File > 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.

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 Volume Replicator replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See "Tasks for a new disaster recovery installation—additional applications" on page 324.

Section 6

Disaster Recovery

- Chapter 13. Disaster recovery: Overview
- Chapter 14. Deploying disaster recovery: New application installation
- Chapter 15. Testing fault readiness by running a fire drill

Chapter 13

Disaster recovery: Overview

This chapter includes the following topics:

- About a disaster recovery solution
- Need for implementing a disaster recovery solution
- Overview of the recovery process
- Components of Volume Replicator that enable disaster recovery

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 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 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 Volume Replicator or hardware replication.

For details on configuring SFW HA disaster recovery using the wizard:

See "Tasks for deploying InfoScale Storage and Volume Replicator with Microsoft failover clustering" on page 522.

For details on SFW and Volume Replicator configuration with Microsoft clustering:

See "Tasks for deploying InfoScale Storage and Volume Replicator with Microsoft failover clustering" on page 522.

The following illustration shows the SFW HA-Volume Replicator configuration with VCS. The example has one disk group on each site for the application. Note that a Volume Replicator 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-Volume Replicator configuration with VCS

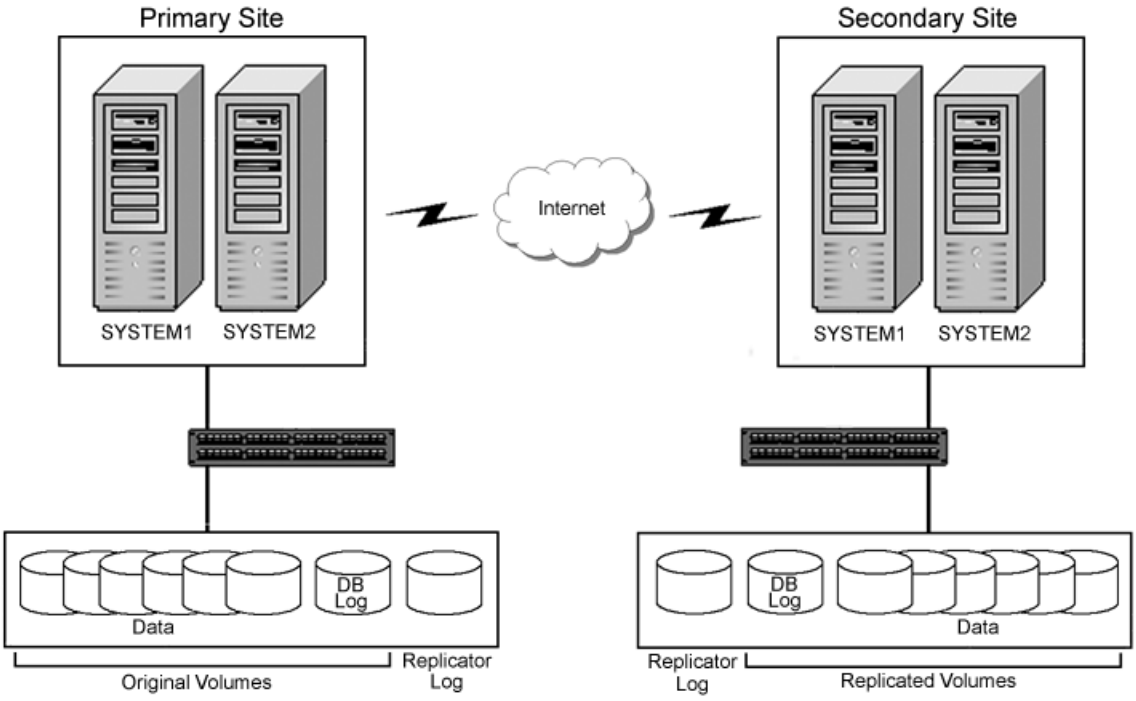

Need for implementing a disaster recovery solution

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 previous 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 Volume Replicator highly available.

Figure 13-2 Typical disaster recovery configuration setup

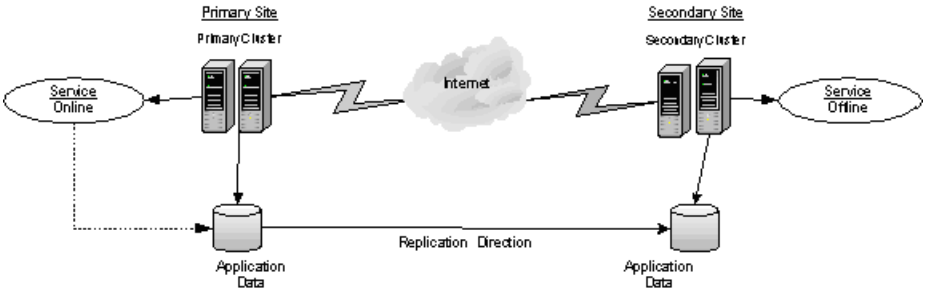

If a disaster, such as an earthquake, causes a failure at the primary site, a host on the secondary site can take over the role of the primary host to make the data accessible and restore the application services and data to users.

Figure 13-3 Recovery situation after a disaster occurs

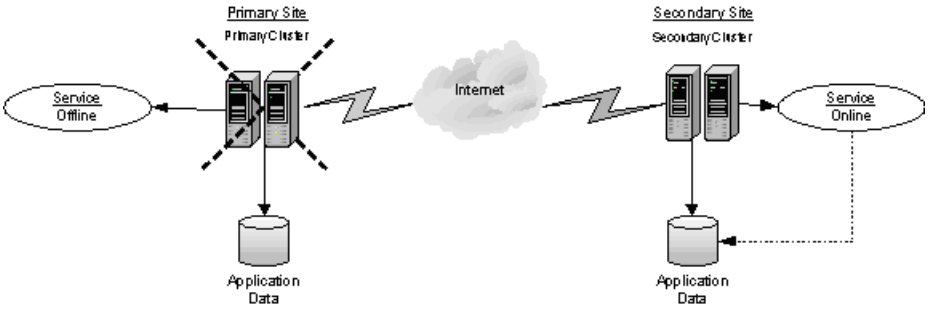

Components of Volume Replicator that enable disaster recovery

This topic provides information about components of Volume Replicator 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.

Volume Replicator (Volume Replicator) 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. Volume Replicator works as an integrated component of Storage Foundation. Any application, even with existing data, can be configured to use Volume Replicator transparently. For more information on Volume Replicator, refer to the *Volume Replicator Administrator's Guide*.

Modes of replication

Volume Replicator 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 *Volume Replicator Administrator's Guide*.

Features of Volume Replicator that help in disaster recovery

While many of the components described above are replicated at the disaster recovery site through conventional means, Volume Replicator solves the difficult problem of replicating the user database.

Refer to the following information on how Volume Replicator helps with disaster recovery in any application environment:

- Write Order Fidelity:
  Volume Replicator 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:
  Volume Replicator 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:
  Volume Replicator 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 Volume Replicator 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:

Volume Replicator provides a replication technology that works with heterogeneous storage hardware. Volume Replicator 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.

Chapter 14

Deploying disaster recovery: New application installation

This chapter includes the following topics:

- Tasks for a new disaster recovery installation—additional applications
- Tasks for setting up DR in a non-shared storage environment
- Notes and recommendations for cluster and application configuration
- Reviewing the configuration
- Configuring the storage hardware and network
- About managing disk groups and volumes
- Setting up the secondary site: Configuring the InfoScale product and setting up a cluster
- Verifying that your application or server role is configured for HA at the primary site
- Setting up your replication environment
- Assigning user privileges (secure clusters only)
- About configuring disaster recovery with the DR wizard
- Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)

- Creating temporary storage on the secondary site using the DR wizard (array-based replication)
- Installing and configuring the application or server role (secondary site)
- Cloning the service group configuration from the primary site to the secondary site
- Configuring the application service group in a non-shared storage environment
- Configuring replication and global clustering
- Creating the replicated data sets (RDS) for Volume Replicator replication
- Creating the Volume Replicator RVG service group for replication
- Configuring the global cluster option for wide-area failover
- Verifying the disaster recovery configuration
- Establishing secure communication within the global cluster (optional)
- Adding multiple DR sites (optional)
- Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment
- Maintaining: Normal operations and recovery procedures (Volume Replicator environment)
- Recovery procedures for service group dependencies

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 “Tasks for a new high availability (HA) installation—additional applications” on page 62.

You can also configure disaster recovery for a primary site that is configured as a replicated data cluster.

See “Tasks for a new replicated data cluster installation—additional applications” on page 236.

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 Volume Replicator 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. Veritas recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See "About the Solutions Configuration Center" on page 23.

Note: If you are using non-shared storage (dynamic disk groups monitored using VMNSDg agent), you cannot use the DR wizard to configure disaster recovery. You have to set up DR manually. Refer to the separate workflow table available for configuring DR manually:

See Table 14-2 on page 328.

This chapter describes the process for any generic application or applications such as FileShare, PrintShare, and IIS.

The following table outlines the high-level objectives and the tasks to complete each objective.

Table 14-1 Task list for deploying disaster recovery

Objective	Tasks
See "Reviewing the configuration" on page 333.	Understanding active-passive configuration and site failover in a DR environment
See "Configuring the storage hardware and network" on page 335.	■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which the application will be installed

Table 14-1 Task list for deploying disaster recovery *(continued)*

Objective	Tasks
See "Setting up the secondary site: Configuring the InfoScale product and setting up a cluster" on page 338.	■ Installing InfoScale Enterprise Refer to the *Veritas InfoScale Installation and Upgrade Guide*. ■ Configuring the cluster using the Cluster Server Configuration Wizard
See "Verifying that your application or server role is configured for HA at the primary site" on page 357.	Verifying that the application has been configured for high availability at the primary site
See "Setting up your replication environment" on page 357.	Ensuring replication prerequisites for your selected method of replication are met before running the DR wizard
See "Assigning user privileges (secure clusters only)" on page 364.	For secure clusters only, assigning user privileges
See "Configuring disaster recovery with the DR wizard" on page 367.	■ 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
See "Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)" on page 370.	(Volume Replicator replication option) Cloning the storage configuration on the secondary
See "Creating temporary storage on the secondary site using the DR wizard (array-based replication)" on page 374.	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Using the DR wizard to create temporary storage for installation on the secondary site
See "Installing and configuring the application or server role (secondary site)" on page 379.	■ Reviewing the prerequisite checklist ■ Installing the application
See "Cloning the service group configuration from the primary site to the secondary site" on page 381.	Cloning the service group configuration from the primary to the secondary site using the DR wizard

Table 14-1 Task list for deploying disaster recovery *(continued)*

Objective	Tasks
See "Configuring replication and global clustering" on page 385.	■ (Volume Replicator 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
See "Verifying the disaster recovery configuration" on page 410.	Verifying that the secondary site has been fully configured for disaster recovery
See "Establishing secure communication within the global cluster (optional)" on page 412.	Adding secure communication between local clusters within the global cluster (optional task)
See "Adding multiple DR sites (optional)" on page 414.	Optionally, adding additional DR sites to a Volume Replicator environment
See "Possible task after creating the DR environment: Adding a new failover node to a Volume Replicator environment" on page 414.	Completing required tasks when adding a new failover system to either the primary or secondary site in a Volume Replicator environment
See "Maintaining: Normal operations and recovery procedures (Volume Replicator environment)" on page 417.	■ Monitor replication ■ Perform planned migration ■ Complete the recovery procedures after the primary site goes down

Tasks for setting up DR in a non-shared storage environment

The following table outlines the high-level objectives and tasks for a creating a single-node DR configuration at the secondary site. Refer to this table if you are setting up DR in a non-shared storage environment (dynamic disk groups configured using VMNSDg agent).

You cannot use the DR wizard to configure disaster recovery in a non-shared storage environment. You have to configure DR manually.

Note: Some procedures (for example, configuring Volume Replicator replication) are common if you are setting up a DR or an RDC configuration. To avoid duplication, the topics referenced in this table point to the procedures described in the RDC chapter covered earlier.

Table 14-2 Non-shared storage: Configuring Disaster Recovery

Objectives	Tasks
Install InfoScale Enterprise and configure the cluster on the secondary site.	■ Verify the software and hardware prerequisites. ■ Set up the network and storage. ■ Install InfoScale Enterprise. **Caution:** Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. ■ Configure disk groups and volumes. See "Setting up the secondary site: Configuring the InfoScale product and setting up a cluster" on page 338.
Verify that the application has been configured for high availability at the primary site.	Verify that the application has been configured for high availability at the primary site and that the service groups are online. See "Verifying that your application or server role is configured for HA at the primary site" on page 357.
Set up the replication prerequisites.	■ Ensure that Volume Replicator replication prerequisites are met. ■ Configure the VxSAS service for Volume Replicator, specifying the cluster nodes at both primary and secondary sites. See "Setting up security for Volume Replicator" on page 531.
(Secure cluster only) Assign user privileges.	For a secure cluster only, assign user privileges. See "Assigning user privileges (secure clusters only)" on page 364.
Install the application at the secondary site.	■ Review. the prerequisite checklist ■ Install the application. See "Installing and configuring the application or server role (secondary site)" on page 379.

Table 14-2 Non-shared storage: Configuring Disaster Recovery *(continued)*

Objectives	Tasks
Configure the application service group for VCS (secondary site).	■ Configure the application service group manually using the Cluster Manager (Java Console). ■ Ensure that the name of the service group is the same as that on the primary site. See "Configuring the application service group in a non-shared storage environment" on page 385.
Set up the replicated data sets (RDS) for Volume Replicator replication.	■ Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites. ■ Use the Setup Replicated Data Set Wizard to create Replicator Log volumes for the primary and secondary sites. See "Creating the replicated data sets (RDS) for Volume Replicator replication" on page 404.
Create the Volume Replicator RVG service group. (Repeat the steps separately on primary and secondary sites.)	■ Create the Volume Replicator RVG service group for the replicated volume group. ■ Use the Cluster Manager (Java Console) to manually create the service group. ■ Create the RVG service group at the primary site and the secondary site separately. ■ Bring the RVG service group online on the primary site. See "Creating the Volume Replicator RVG service group for replication" on page 404.
Configure the global cluster option for wide-area failover.	■ Link clusters (adding a remote cluster to a local cluster). ■ Convert the application service group that is common to all the clusters to a global service group. ■ Convert the local service group to a global group. ■ Bring the global service group online. See "Configuring the global cluster option for wide-area failover" on page 405.
Verify the disaster recover configuration.	Verify that the secondary site has been fully configured for disaster recovery. See "Verifying the disaster recovery configuration" on page 410.

Table 14-2 Non-shared storage: Configuring Disaster Recovery *(continued)*

Objectives	Tasks
(Optional) Add secure communication.	Add secure communication between local clusters within the global cluster (optional task). See "Establishing secure communication within the global cluster (optional)" on page 412.
(Optional) Add additional DR sites.	Optionally, add additional DR sites to a Volume Replicator environment. See "Adding multiple DR sites (optional)" on page 414.
Handlr service group dependencies after failover.	If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site. See "Recovery procedures for service group dependencies" on page 420.

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
  https://sort.veritas.com/documents

  Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

  The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

  Refer to the following Microsoft knowledge database article for more details:

  http://technet.microsoft.com/en-us/library/dd184075.aspx

- 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 *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. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

  Static IP addresses are required for the following purposes:

  - One static IP address 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 Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
  - For Volume Replicator 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 SFW HA 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 *Cluster Server Bundled Agents Reference Guide*.

- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- 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.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.
- For a Replicated Data Cluster, install only in a single domain.
- 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)
- This is applicable for a Replicated Data Cluster configuration.
  This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, 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.
- To configure a RDC cluster, you need to 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.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	The following types of IPv6 addresses are supported: ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	LLT over UDP is supported on both IPv4 and IPv6. You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.
VCS agents, wizards, and other components	VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above). The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

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 "About a disaster recovery solution" on page 316.

Figure 14-1 DR configuration primary site

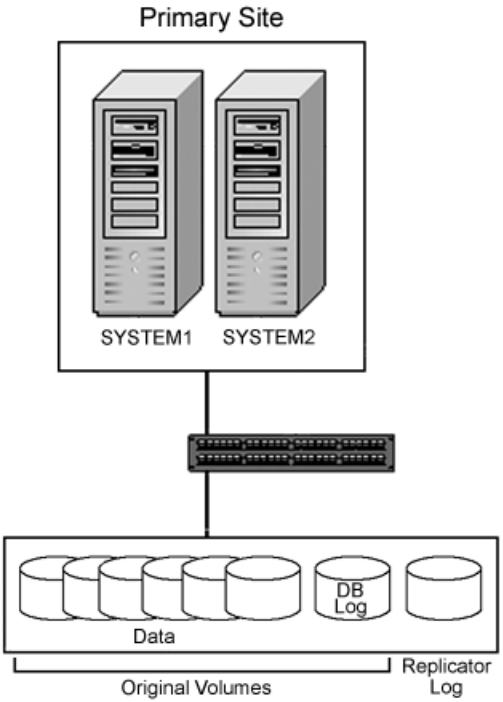

Supported disaster recovery configurations for service group dependencies

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 the *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.

In a hardware replication environment, the Disaster Recovery wizard supports one level of dependency (one child). 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.

In a Volume Replicator environment, the wizard cannot configure DR for a service group that has a child and you will need to configure the secondary site manually. For more information on configuring Volume Replicator, see the *Volume Replicator Administrator's Guide*. For more information on configuring GCO, see the *Cluster Server Administrator's Guide*.

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

1. Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
2. Connect the network adapters on each system.
   - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
   - Veritas 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.

To verify the DNS settings and binding order for all systems

1. Open the Control Panel by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.

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, click **Adapter settings**.

4. Ensure the public network adapter is the first bound adapter by following these steps sequentially:

   - In the Network Connections window, click **Advanced > 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)** or **Internet Protocol Version 6 (TCP/IPv6)** check box, depending on which protocol your network is using.
   - 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 that the **Register this connection's address in DNS** check box is selected.

12 Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.

13 Click **OK**.

About managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a cluster dynamic 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.

Managing disk groups and volumes involves the following:

- See "Importing a disk group and mounting a volume" on page 337.
- See "Unmounting a volume and deporting a disk group" on page 338.

Note: (Disaster recovery configurations only) 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 *Storage Foundation Administrator's Guide* for more information.

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 the InfoScale product and setting up a cluster

Begin with verifying that the requirements are met on the secondary site.

See "Notes and recommendations for cluster and application configuration" on page 330.

Then, continue with the following:

- See "About installing the Veritas InfoScale products" on page 503.
- See "Configuring the cluster using the Cluster Configuration Wizard" on page 345.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Installing the server components using the installation wizard

The product installation wizard allows you to install the product on multiple systems at a time.

Before you begin to install the product ensure that you have reviewed the installation prerequisites, licensing, and the product co-existence details.

Note: If you plan to install InfoScale Storage in an active Microsoft Failover Cluster, ensure that you have reviewed the applicable pre-requisites, and use the "rolling-install" procedure to perform the product installation. To use the "rolling-install" procedure, install InfoScale Storage first on the inactive cluster node. Then move the cluster resources to the other node and install the product on the now inactive node.

Perform the following steps to install the server components

1. Download the installation package from the following location:

   https://sort.veritas.com

2. Allow the autorun feature to start the installation or double-click **Setup.exe**.

   The CD browser appears.

   Note: If you install the software using the product software disc, the CD browser displays the installation options for all the products that are specified earlier. However, if you download the installation package from the Veritas website, the CD browser displays the installation options only for the product to be installed.

   ---

3 Click the required product-specific tab and then click the link to install the components.

Note: The client components are installed by default along with the server components. However, the client components are not installed if the system is a server core machine.

In addition to the product-specific tabs, the CD browser also provides the following links:

Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
Windows Data Collector	Click to verify that your configuration meets all pertinent software and hardware requirements.
SORT	Click to access the Veritas Services and Operations Readiness Tools (SORT) site. In addition to the product download you can also download the custom reports about your computer and Veritas enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
Browse Contents	Click to view the software disc contents.
Technical Support	Click to contact Veritas Technical Support.

4 On the Welcome panel, review the list of prerequisites and click **Next**.

Note that the **Check for product updates** check box is selected by default. The wizard searches for the available product updates on the SORT website. You can then download and apply the available updates. If you do not want to apply the available updates, clear the selection of **Check for product updates** check box.

5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.

The **Participate in the Veritas Product Improvement Program by submitting system and usage information anonymously** check box is selected by default. The Product Improvement Program allows the wizard to collect installation, deployment, and usage data and submit it anonymously to Veritas. The collected information helps identify how customers deploy and use the product. If you do not want to participate in the Product Improvement Program, clear the selection of the check box.

6 On the Product Updates panel, review the list of available product updates.

This panel appears only if you have selected the **Check for product updates** check box on the Welcome panel.

The product updates comprise of the pre-installation patches, post-installation patches, high availability Agents, and array-specific modules. The panel lists the available pre-installation patches and the post-installation patches. Download and apply the pre-installation patches in the sequence shown in the table and rerun the wizard. After the successful installation of the product, apply the post-installation patches. Also download and install the high availability agents and array-specific modules from the SORT website.

7 On the System Selection panel, select the systems and the desired Installation and Product options:

You can select the systems in one of the following ways:

- In the System Name or IP text box, manually type the system name or its IP address and click **Add**.

  Note: The wizard does not support the Internet Protocol version 6. To add the systems having Internet protocol version 6, you must type the system name.

  The local host is populated by default.

- Alternatively, browse to select the systems.
  The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks, and notes the details in the Verification Details box. To review the details, select the desired system.

To select the installation and product options, perform the following tasks on each of the selected system.

Note: To apply the selection to multiple systems, select the system for which you have selected the installation and product options and then click **Apply to multiple systems**.

See "Applying the selected installation and product options to multiple systems" on page 470.

- By default the wizard uses %ProgramFiles%\Veritas as the installation directory. To customize the installation directory, click **Browse** and select the desired location. Click **OK**.

  Install the product at the same location on all the systems.

  If you are upgrading the product, the installation directory is selected by default.

  Note: The installation directory must contain only English characters, if:

  - You plan to configure the cluster for single sign-on authentication.

  - Your system runs a non-English locale operating system.

- Select the required license type from the **License key** drop-down list.

  Note: The default license type is "Keyless".

  If you select the "Keyless" license type, all the available product options are displayed and are selected by default.

  If you select "User entered license key" as your license type, the License Details panel appears by default. On the License Details panel, enter the license key and then click **Add**. You can add multiple licenses for the various product options you want to use.

  The wizard validates the entered license keys and displays the relevant error if the validation fails. After the validation is complete, click **OK**.

- From the list of product options, select the options to be installed.

  The options differ depending on the product you install.

8 On the System Selection panel, click **Next**.

Note that the wizard fails to proceed with the installation, unless all the selected systems have passed the validation checks and are ready for installation. In case the validation checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the installation, select the system and click **Re-verify** to re-initiate the validation checks for this system.

9 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This selection reboots all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

10 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

If an installation is not successful on any of the systems, the status screen shows a failed installation.

Note: During the upgrade, the Installation panel displays a list of services and processes running on the systems. Select a system to view the services and processes running on it and review the list.

The wizard stops the product-specific services and discovers the processes running, if any, on the systems. These processes need to be stopped to proceed with the operation. Click **Next** to forcefully stop the processes and proceed with the operation. Alternatively, you can manually stop the processes. If the services or processes cannot be stopped, the operation fails. Rectify the error and then click **Retry** to validate the affected system again. Click **Retry All** to validate all the systems again.

In case you want to proceed with the upgrade without stopping a particular process, contact Veritas Technical Support.

---

11 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed on any of the system, refer to the log file for details. You may have to re-install the software.

12 On the Finish panel, click **Finish**.

If you had chosen to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot these systems.

This completes the product installation. You must now proceed to configure the required components. Refer to the component-specific guides for more details about the configuration tasks.

Note: If you have installed InfoScale Storage with Microsoft Failover Cluster, but if Microsoft failover cluster is not yet configured, you must register the InfoScale Storage resources, after configuring the Microsoft failover cluster software.

See "Registering the InfoScale Storage resource DLLs" on page 471.

However, if you have installed InfoScale Storage in an active Microsoft Failover Cluster, then you must remove the physical disk resources for all the basic disks. You must do so before configuring the SFW cluster disk groups. Failing this, a reservation conflict occurs.

Applying the selected installation and product options to multiple systems

To apply the selected installation and product options to multiple systems, perform the following steps:

1. Click on any one of the selected systems and select the desired installation and product options.
2. Click **Apply to multiple systems**.
3. On the Apply Installation Options panel, select the installation options to be applied and then select the desired systems. Click **OK**.

   Note: The installation directory is selected by default on the systems where the product is being upgraded. The selected **Install Directory** option does not apply to these systems.

Installing the client components

Note: Client components cannot be installed on server core systems.

Perform the following steps to install the client components

1. Log on to any node in the cluster.
2. From the Veritas Services and Operations Readiness Tools (SORT) site, download the client-specific .msi file.

   https://sort.veritas.com/

msi name	Client component name
vrtsvcscscm.msi	VCS Cluster Manager (Java Console)
VRTSobgui.msi	Veritas Enterprise Administrator (VEA)

3. Follow the instructions that the install program provides, to complete the installation.

Configuring the cluster using the Cluster Configuration Wizard

After installing the software, set up the components required to run Cluster Server. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, the user account for the VCS Helper service, and provides an option for configuring the VCS Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for notification and global clusters (GCO). You can also use VCW to modify or delete cluster configurations.

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 VCW to remove the node from the cluster, rename the system, and then run VCW again to add that system to the cluster.

Note the following prerequisites before you proceed:

- The required network adapters, and SCSI controllers are installed and connected to each system.

  To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.
- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.

- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

  Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list.
  When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
  Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.

- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

1. Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
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.

   To discover information about all systems and users in the domain, do the following:

   - Clear **Specify systems and users manually**.
   - Click **Next**.
     Proceed to step 8.

   To specify systems and user names manually (recommended for large domains), do the following:

   - Select **Specify systems and users manually**.
     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. 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.

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**.

   Select the 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.

   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.
   - Product is either not installed or there is a version mismatch.

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

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

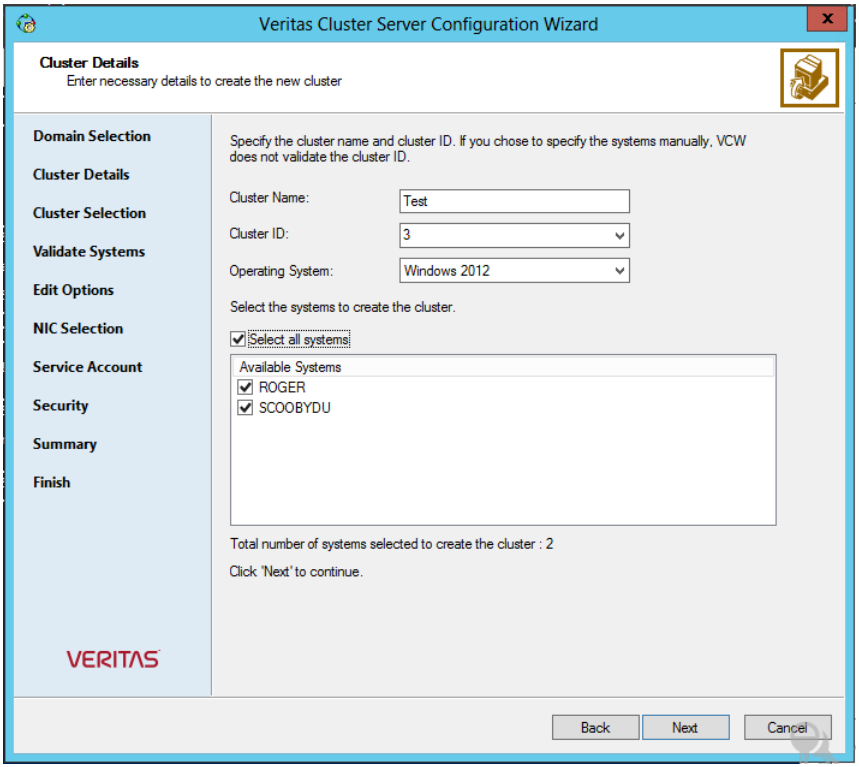

Specify the cluster details as follows:

Cluster Name	Type a name for the new cluster. Veritas 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 65535. **Note:** 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. All the systems in the cluster must have the same operating system and architecture. For example, you cannot configure a Windows Server 2008 R2 system and a Windows Server 2012 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 step 9, proceed to the next step. Otherwise, proceed to step 12.

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 using IPv4 or IPv6 network.

Do one of the following:

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

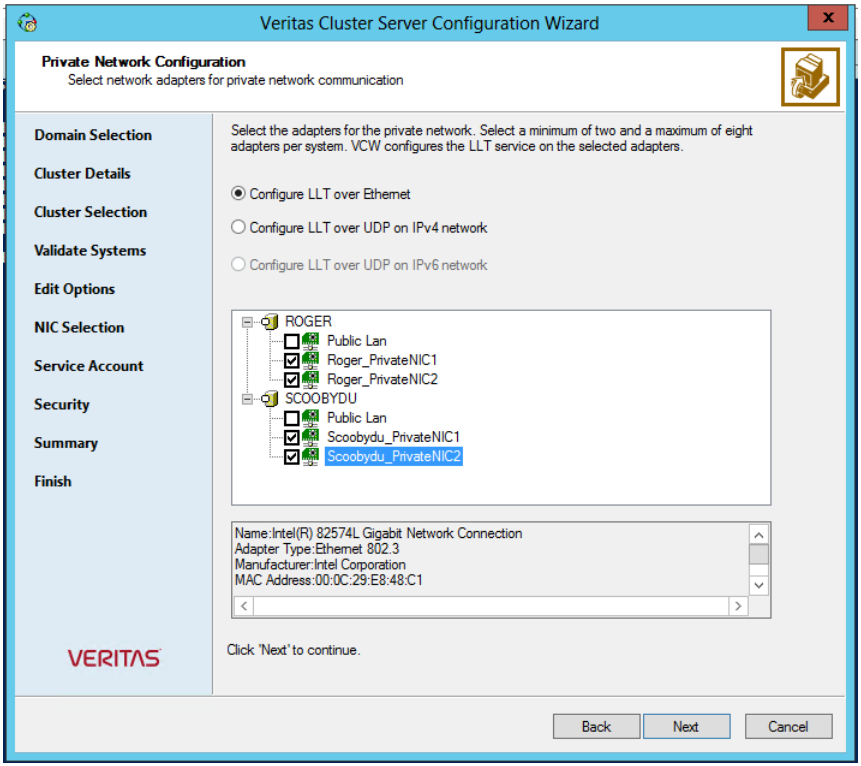

- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.
  Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends 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. Veritas 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 on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.
  The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the VCS private network.
- 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. In case of IPv4, each IP address can be in a different subnet.
  The IP address is used for the VCS private communication over the specified UDP port.
- 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, 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 are used for the private network communication.

12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service.

The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

Specify the domain user details 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, 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 communications and then click **Next**.

Do one of the following:

- To use VCS cluster user privileges, 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 Cluster Manager (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.

  Veritas recommends that you specify a new user name and password.

- To use the single sign-on feature, click **Use Single Sign-on**.

  In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.

  The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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 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.

16 On the Cluster Service Components panel, select the components to be configured in the ClusterService group and then click **Next**.

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 354.
- 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.

  Note: Configure GCO using this wizard only if you are configuring for disaster recovery but not using the DR wizard.

  Alternatively, consider configuring GCO using the DR wizard.

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 then 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 then click **Next**.

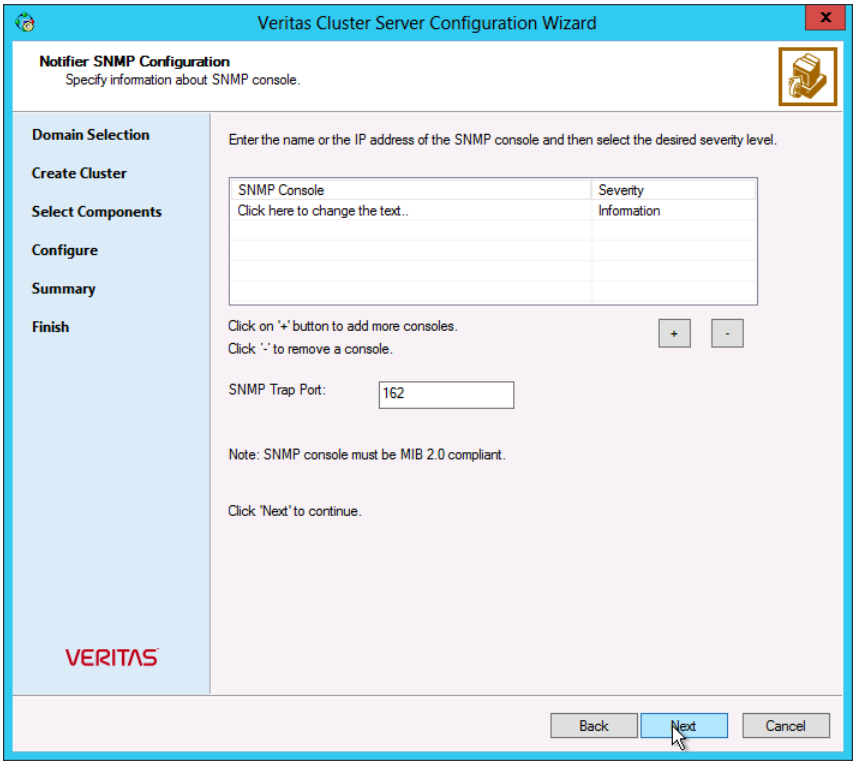

Do the following:

- 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 the **+** icon to add a field; click the **-** icon 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 then click **Next**.

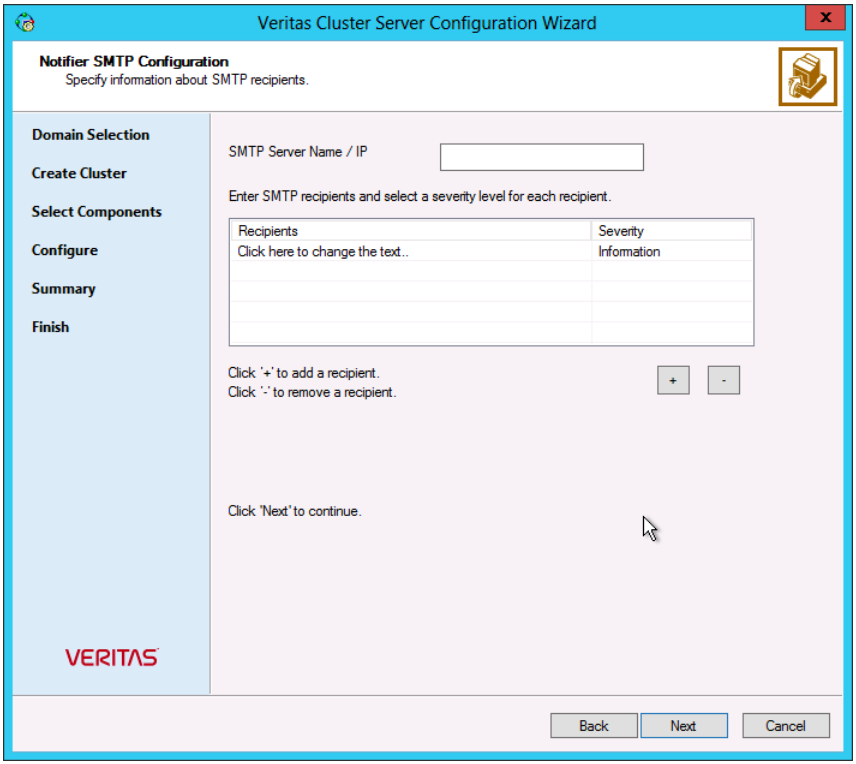

Do the following:

- 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 the **+** icon to add fields; click the **-** icon to remove a field.

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

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that 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 starts and click **Configure**.

6 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 "Tasks for a new high availability (HA) installation—additional applications" on page 62.

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 "Tasks for a new replicated data cluster installation—additional applications" on page 236.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Volume Replicator (Volume Replicator)
- EMC SRDF

- Hitachi TrueCopy

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

For array-based hardware replication, you can use any replication agent supported by 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 401.

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:

- See "Setting up security for Volume Replicator" on page 531.
- See "Requirements for EMC SRDF array-based hardware replication" on page 360.
- See "Requirements for Hitachi TrueCopy array-based hardware replication" on page 362.

Setting up security for Volume Replicator

If you use Volume Replicator for replication, you must configure the Veritas Volume Replicator Security Service (VxSAS) on all the cluster nodes.

In a Replicated Data Cluster environment, you must configure the service on all the nodes in the primary zone as well as the secondary zone.

For details on this required service, see the *Volume Replicator Administrator's Guide*.

After you install InfoScale Storage or InfoScale Enterprise, launch the Veritas Volume Replicator Security Service Configuration Wizard. This wizard lets you complete the Volume Replicator security service configuration.

To do so, launch the wizard after you install InfoScale Enterprise on both the primary and secondary nodes. Then, when you run the wizard, you can specify the primary and secondary sites in one step.

Prerequisites for configuring VxSAS

- The wizard requires you to be logged on with administrative privileges.

- The account that 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.
- The systems on which you want to configure VxSAS must be accessible from the local system.

To configure VxSAS

1. Launch the Veritas Volume replicator Security Service Configuration Wizard from **Start > All Programs > Veritas > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   Optionally, run `vxsascfg.exe` from the command prompt of the required machine.

2. Read the information provided on the Welcome page and click **Next**.

3. Complete the Account Information panel as follows:

Account name (domain\account)	Enter the administrative account name.
Password	Specify a password

   If you have already configured VxSAS for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring VxSAS on the other hosts.

   Click **Next**.

4. On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 lets you specify the domain name. Click **Add** to add the name to the Selected domains list.

   Click **Next**.

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 VxSAS.

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.

When configuring VxSAS in a firewall setup, the VxSAS wizard may not be able to configure the machines that are across the firewall, although the Host Selection dialog may list these nodes. In this case, configure VxSAS locally on the machines that are across the firewall.

Click **Back** to change any information you had provided earlier.

7 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 the *Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide*.

Before using the DR wizard, review the following:

- See "Software requirements for configuring EMC SRDF" on page 361.
- See "Replication requirements for EMC SRDF" on page 361.

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.

Note: In addition, the agent requires that the 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 the *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

Before using the DR wizard, review the following:

- See "Software requirements for Hitachi TrueCopy" on page 362.
- See "Replication requirements for Hitachi TrueCopy" on page 363.

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

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:

  `systemDriver\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 *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
```

About configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (Volume Replicator replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure Volume Replicator 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:

- InfoScale Enterprise 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.
- 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.
- For IPv4 networks, static IP addresses are available to enter for the following (for IPv6, they are generated during configuration):
  - One static IP address per application service group to be cloned.
  - One static IP address at each site for configuring GCO.
  - If using Volume Replicator for replication, a minimum of one static IP address per site for each application instance running in the cluster.
- The service group to be cloned can use either IPv4 IP addresses or IPv6 addresses but not a mixture of both.
- To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed.
- For Volume Replicator replication, the service group to be cloned cannot contain a child service group.
- 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 Volume Replicator configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- See "Setting up security for Volume Replicator" on page 531.
- See "Requirements for EMC SRDF array-based hardware replication" on page 360.
- See "Requirements for Hitachi TrueCopy array-based hardware replication" on page 362.

Configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

This procedure describes how to configure disaster recovery using the wizard.

To start configuring disaster recovery with the DR wizard

1. Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**, or, on Windows 2012 operating systems, from the **Apps** menu.
2. 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.

3. In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.

4 In the System Selection panel, provide information in the **System Name** field:

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 **localhost** to connect to the system.

Click **Next**.

5 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.

For a hardware replication environment, 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 Volume Replicator environment, the DR wizard does not support configuring DR for a service group that has a child. If you select a service group that has a child, you will receive an error message when you select the Volume Replicator replication method later in the wizard.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

6 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**.

7 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 Volume Replicator (Volume Replicator) and the Global Cluster Option (GCO)	Select this option if you want to configure Volume Replicator replication. Select this option even if you plan to configure Volume Replicator replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a Volume Replicator 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 Volume Replicator option, the wizard will warn you that you cannot use Volume Replicator replication for the disaster recovery site.
Configure EMC SRDF and the Global Cluster Option (GCO)	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. 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.
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	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. Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.

Configure the Global Cluster Option (GCO) only	If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option. 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. 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 Volume Replicator replication environment.

Click **Next**.

8 Continue with the next DR configuration task.

For Volume Replicator replication:

See "Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)" on page 370.

For array-based replication:

See "Creating temporary storage on the secondary site using the DR wizard (array-based replication)" on page 374.

Cloning the storage on the secondary site using the DR wizard (Volume Replicator 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.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you have not yet started the wizard, refer to the following topic before continuing with the storage cloning procedure:

See "Configuring disaster recovery with the DR wizard" on page 367.

To clone the storage configuration from the primary site to the secondary site (Volume Replicator 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 Volume Replicator 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	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. ■ 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 **>>** button 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 **>>** button 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.

Note: On the VEA GUI of the secondary site, a Windows dialog box might appear prompting you to format a disk. Click **Cancel** to close the dialog.

The appearance of this dialog box has no impact on the operations being performed by the DR wizard. You can safely ignore it.

---

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.

  If the system must be restarted once application installation is complete, and you are running the wizard from a local node, 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.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

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, refer to the following topic before continuing with the storage cloning procedure:

See "About configuring disaster recovery with the DR wizard" on page 365.

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 The default path is `C:\HORCM\etc`, where `C` is the system drive.
HORCM files location	Path to the horcm configuration files (horcmnn.conf) The default path is: `C:\Windows`, where `C` is the system drive. The `horcm` configuration file is required by the RAID Manager on all nodes; however, the wizard does not validate its presence.

4. 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.

5 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.
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**.

6 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**.

7 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**.

8 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**.

Note: If SCSI-3 support is enabled for using Persistent Group Reservations (PGR), and if one of the selected disks is not SCSI-3 compliant, the following error is displayed: “Unable to reserve a majority of dynamic disk group members. Failed to start SCSI reservation thread.”

Recommended action: Click **Finish** to exit the wizard. Either replace the non-compliant disk with a SCSI-3 compliant disk, or enable SCSI-2 support, and then run the wizard again.

---

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)

Note: The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems.

- See "Installing the FileShare application" on page 379.
- See "Installing the PrintShare application" on page 379.
- See "Installing the IIS application" on page 380.
- See "Installing additional applications" on page 380.

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 (VMDg) or non-shared storage (VMNSDg).
- When installing and configuring a new file server shared directory, create the disk groups and volumes on the shared storage (VMDg) or non-shared storage (VMNSDg) and subsequently create the directory structure for the file shares.
- If your configuration already has a file server with shares on the local storage, then move these shares to the shared storage (VMDg) or non-shared storage (VMNSDg) using practices recommended by Microsoft.

Installing the PrintShare application

Points to note when installing PrintShare:

- Make sure that 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 the *Cluster Server Administrator's Guide*.
- Make sure that the printer is connected to the network and is configured with an IP address.

Note: The agent for PrintShare is not supported on systems that run Windows Server 2012 or Windows Server 2012 R2. Therefore, you will not be able to configure print shares on such systems.

Installing the IIS application

Points to note when installing IIS:

- Verify that IIS is installed and configured identically on all nodes hosting the service group. Verify that the sites to be monitored are on the shared storage (VMDg) or non-shared storage (VMNSDg).
- Import the 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 (VMDg) or non-shared storage (VMNSDg) and place the site content in that folder.
- Change the default home directory path for all the IIS sites to be monitored to a location on the shared storage (VMDg) or non-shared storage (VMNSDg). See the IIS documentation for instructions.
- For existing websites, stop the sites and then move the website content to volumes on the shared storage (VMDg) or non-shared storage (VMNSDg). 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 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 (VMDg) or non-shared storage (VMNSDg).

Cloning the service group configuration from the primary site to the secondary site

Ensure that the SQL Server Full-Text Search service on the secondary site is configured to start in the manual mode and is initially in the stopped state.

Before cloning the service group on the secondary site, verify that you have installed the application on the secondary site.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you are launching the wizard for the first time, refer to the following topic for additional information:

See "About configuring disaster recovery with the DR wizard" on page 365.

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 by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.

   Expand the Solutions for Solutions for Additional Applications tab.

   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 Volume Replicator 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	For an IPv4 network, 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 IPv6, select the network from the dropdown list. An IP address will be generated from the network. For the MACAddress attribute select the appropriate public NIC from the drop-down list. For IPv6 available NICs are those belonging to the selected IPv6 network.

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 the application service group in a non-shared storage environment

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to clone the application service group created at the primary site if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the service group manually using the Cluster Manager (Java Console).

Note the following before configuring the service group at the secondary site:

- Ensure that the application agent resources, the Lanman resource (if configured), and the IP resource is offline in the service group on the primary site. The remaining resources, including the storage resources, must be online.
- Ensure that the name of the service group is the same as that on the primary site.
- After configuring the service group do not bring it online on the secondary site at this time. You can bring it online later after completing all the DR configuration steps.
  See "Configuring the service group in a non-shared storage environment" on page 163.

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 Volume Replicator or an agent-supported array-based hardware replication.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

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:

- See "Configuring Volume Replicator replication and global clustering" on page 386.
- See "Configuring EMC SRDF replication and global clustering" on page 394.
- See "Configuring Hitachi TrueCopy replication and global clustering" on page 397.
- See "Configuring global clustering only" on page 401.

Configuring Volume Replicator replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure Volume Replicator replication and global clustering.

Before you begin, ensure that you have met the following prerequisites:

- Ensure that Volume Replicator Security Service (VxSAS) is configured at the primary and secondary site.
  See "Setting up security for Volume Replicator" on page 531.
- Ensure that you have set the appropriate IP preference, whether Volume Replicator should use IPv4 or IPv6 addresses. The default setting is IPv4.
  When you specify host names while configuring replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP protocol Volume Replicator uses to resolve the host names.
  Use Veritas Enterprise Administrator (VEA) (**Control Panel > VVR Configuration > IP Settings** tab) to set the IP preference.
- 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 Volume Replicator replication and global clustering with the DR wizard.

To configure Volume Replicator 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 DR 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 by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
   - Expand the Solutions for Additional Applications tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
3. On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
4. On the Replication Methods panel, click **Configure Volume Replicator and the Global Cluster Option (GCO)**. Click **Next**.
5. In the Internet Protocol panel, select IPv4 or IPv6 depending on which type of network you are using. (You must use the same on primary and secondary sites.) Click **Next**.
6. 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.
7. 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 **>** button 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 **<** button to move the volumes into the Available Volumes pane.
Primary SRL	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. Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.
Secondary SRL	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. Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.
Add RVG	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, you can use the **Add RVG** option to organize them differently, based on your specific requirements.
Delete RVG	Click this option to delete any of the existing RVGs related to the DR set up that you are creating.
Start Replication after the wizard completes	Select this check box to start replication automatically after the wizard completes the necessary configurations. Once replication is configured and running, deselecting the checkbox does not stop replication.

Click **Advanced Settings** to specify some additional replication properties.

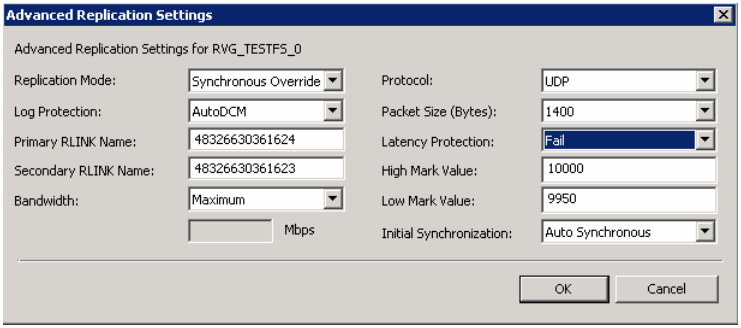

The options on the dialog box are described column-wise, from left to right:

Replication Mode — Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override** (default).

Log Protection — Select the appropriate log protection from the list:

- **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 **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 — 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	By default, Volume Replicator replication uses the maximum available bandwidth. You can select **Specify** to specify a bandwidth limit. The default unit is Megabits per second (Mbps) and the minimum allowed value is 1 Mbps.
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	By default, latency protection is set to **Off**. 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. This **Override** option behaves like the **Off** option when the Secondary is disconnected and behaves like the **Fail** option when the Secondary is connected.
High Mark Value	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. To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.
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	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. When this option is selected, Volume Replicator 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. If you want to use the **Synchronize from Checkpoint** method then you must first create a checkpoint. 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.

To apply changes to advanced settings, click **OK**.

For additional information on Volume Replicator replication options, refer to the *Volume Replicator Administrator's Guide*.

Click **Next**.

8 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	For IPv4 networks, enter replication IPs that will be used for replication, one for the primary site and another for the secondary site. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site. For IPv6, available NICs are those belonging to the selected network.
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**.

9 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	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
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.

10 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.

Otherwise, click **Next** to implement the settings.

11 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**.

12 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 one static address is available per site for configuring GCO.

See "Requirements for EMC SRDF array-based hardware replication" on page 360.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings.

See "Optional settings for EMC SRDF" on page 397.

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 by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
- Expand the Solutions for Additional Applications tab.
  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.

Otherwise, 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 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 the *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:

Table 14-3 Optional settings for EMC SRDF

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 one static address is available per site for configuring GCO.

See "Requirements for Hitachi TrueCopy array-based hardware replication" on page 362.

The wizard configures the required agent settings. It uses defaults for optional settings.

See "Optional settings for HTC" on page 401.

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 by following these steps sequentially:

   - Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.

     Expand the Solutions for Solutions for Additional Applications tab.

     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.
4. In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
5. 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.

---

6 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.

7 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.

8 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.

9 In the Settings Summary panel, review the displayed information.

If you want to change any of the parameters specified for the replication resource settings or the global cluster settings, click **Back**.

Otherwise, click **Next**.

10 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**.

11 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.

12 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 the *Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Table 14-4 Optional settings for HTC

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 you have one static address is 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 by following these steps sequentially:

   - Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
   - Expand the Solutions for Solutions for Additional Applications tab.
     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	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
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.

If you want to change any of the parameters specified, click **Back**.

Otherwise, 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 (**x**) 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.

Creating the replicated data sets (RDS) for Volume Replicator replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure replication using the Setup Replicated Data Set Wizard.

Configuring Volume Replicator involves setting up the replicated data sets (RDS) 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.

See "Creating the Replicated Data Sets with the wizard" on page 539.

Creating the Volume Replicator RVG service group for replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the replication service group manually using the Cluster Manager (Java Console).

Complete the following procedures first on the node in the primary site. Then repeat all the steps on the node in the secondary site. You must follow the order of the procedures as mentioned.

Refer to the following topics:

- See "Configuring a RVG service group for replication" on page 288.

- See "Creating the RVG service group" on page 288.
- See "Configuring the IP and NIC resources" on page 290.
- See "Configuring the VMDg or VMNSDg resources for the disk groups" on page 292.
- See "Adding the Volume Replicator RVG resources for the disk groups" on page 294.
- See "Linking the Volume Replicator RVG resources to establish dependencies" on page 295.
- See "Deleting the VMDg or VMNSDg resource from the application service group" on page 296.
- See "Configuring the RVG Primary resources" on page 297.
- See "Creating the RVG Primary resources" on page 297.
- See "Linking the RVG Primary resources to establish dependencies" on page 298.
- See "Bringing the RVG Primary resources online" on page 298.
- See "Setting a dependency between the service groups" on page 299.

Configuring the global cluster option for wide-area failover

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

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 and perform the following global cluster operations:

- See "Linking clusters: Adding a remote cluster to a local cluster" on page 406.
- See "Converting a local service group to a global service group" on page 407.
- See "Bringing a global service group online" on page 409.

Linking clusters: Adding a remote cluster to a local cluster

This is applicable only if you are setting up DR manually in a non-shared storage environment.

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.

  Veritas InfoScale products do 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 **Edit > Add/Delete Remote Cluster**.

   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 by entering `hahb -display` in the command window.

The state attribute in the output should show "alive". If the state is unknown, then take the ClusterService group offline and bring it online again.

Converting a local service group to a global service group

This is applicable only if you are setting up DR manually in a non-shared storage environment.

To convert a local service group to a global group

1. From Cluster Explorer, click **Edit > Configure Global Groups**.

   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:

Cluster not in secure mode	Follow these steps sequentially: ■ 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	Follow these steps sequentially: ■ 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.

5 Click **Next**, then click **Finish**.

At this point, you must bring the global service group online from Cluster Explorer.

Bringing a global service group online

This is applicable only if you are setting up DR manually in a non-shared storage environment.

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**.

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 Volume Replicator 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 Volume Replicator replication:
  - Ensure Volume Replicator 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 Volume Replicator 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.
- 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 *Cluster Server Hardware Replication Agent for Hitachi TrueCopy 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 Volume Replicator for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using Volume Replicator 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 process typically consists of the following tasks:

  - Starting a Volume Replicator replication checkpoint
  - Performing a block level backup
  - Ending the Volume Replicator replication checkpoint
  - Restoring the block level backup at the DR site
  - Starting replication from the Volume Replicator replication checkpoint

  To learn more about the process of starting replication from a checkpoint, refer to the *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 Volume Replicator-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 all the previous steps 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 > 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 the previous step 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 all the previous steps 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 rootHashInHex]
  ```

  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 use the following commands:

  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 all the previous steps for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. 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 Volume Replicator 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 Volume Replicator 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 InfoScale Enterprise on the new system and then add the system to the cluster.

To install InfoScale Enterprise and add the system to the cluster

1. For installation instructions:

   See "About installing the Veritas InfoScale products" on page 503.

2. Start the Veritas Cluster Server Configuration Wizard from the Solutions Configuration Center.

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or click the shortcut for the Solutions Configuration Center.

   On Windows 2012 operating systems, use the **Apps** menu.

   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 *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 Volume Replicator 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 to add a new node to the system list for the replication service group.

   Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** or, on Windows 2012 operating systems, use the **Apps** menu.

   If necessary, refer to the *Volume Replicator Administrator's Guide* for information on this procedure.
4. Use the Modify service group option of the FileShare, PrintShare, and IIS Configuration Wizard or Application Configuration Wizard.

   Start the appropriate Configuration Wizard from the Solutions Configuration Center. For example, for FileShare click **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center > High Availability Configuration Wizards > FileShare Configuration Wizard** to add the new node to the system list for the respective application service group.

   On Windows 2012 operating systems, use the **Apps** menu.

   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 (See "Preparing the existing DR environment" on page 415.), 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 (Volume Replicator environment)

This section provides tasks during normal operations of the DR solutions and also describes the recovery process.

Refer to the following topics:

- See "Normal operations: Monitoring the status of the replication" on page 417.
- See "Performing planned migration" on page 418.
- See "Disaster recovery procedures" on page 418.

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 *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-failback option to restore the original primary easily once it becomes available again. When performing Takeover with fast-failback, make sure that you do not select the Synchronize Automatically option.
- Perform Takeover without the fast-failback 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 performing 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 sychronization 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 334.

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 Volume Replicator 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 14-5 Online, local, soft dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	■ 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.	■ Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
The parent service group fails	■ 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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ 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 14-6 Online, local, firm dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	■ 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.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Leave the RVG group online at the primary site.
The parent service group fails	■ 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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ 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, hard dependency link.

Table 14-7 Online, local, hard dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	■ 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.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent). Do not take the RVG group offline at the primary site.
The parent service group fails	■ 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.	■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Chapter 15

Testing fault readiness by running a fire drill

This chapter includes the following topics:

- 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
- Re-creating a fire drill configuration that has changed
- Restoring the fire drill system to a prepared state
- Deleting the fire drill configuration
- Considerations for switching over fire drill service groups

About disaster recovery fire drills

A disaster recovery (DR) plan should include regular testing of an environment to ensure that a DR solution is effective and ready if a disaster strikes. 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

Storage Foundation and High Availability Solutions (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 Volume Replicator (Volume Replicator) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

In the Hitachi TrueCopy or EMC SRDF environments, the Fire Drill Wizard supports only the Gold configuration. For the Silver or Bronze configuration, you must manage (create, restore, delete) the fire drill configurations and run the fire drills manually. For further information about the Gold, Silver, and Bronze configurations, refer to the following documents:

Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide

Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide

---

Note: After upgrading to 6.0.1 or later, the existing fire drill service groups will not be usable. In a Hitachi TrueCopy or EMC SRDF environment, you must manually edit the existing fire drill service groups. In a Volume Replicator environment, you must use the Fire Drill Wizard to re-create them. For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.

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 a fire drill, 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, after completing the fire drill testing for a service group, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible.

  See "Restoring the fire drill system to a prepared state" on page 448.

- 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 Volume Replicator environment" on page 425.

See "About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment" on page 428.

About Fire Drill Wizard operations in a Volume Replicator 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 424.

However, the following additional Fire Drill Wizard operations are specific to a Volume Replicator environment.

Preparing the fire drill configuration

In a Volume Replicator environment, when preparing the fire drill configuration, the wizard does the following:

- Replaces the RVGPrimary resources with VVRSnap resources in the fire drill service group
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill
  While running the wizard, you assign one or more disks for the mirrored volumes. Mirror preparation can take some time, so you can exit the wizard after this step is started and let the preparation continue in the background.
- Sets the `offline-local-firm` dependency between the service groups, where the fire drill service group is the parent and the application service group is the child
- Configures the VVRSnap resource by setting the following attributes to the appropriate values:
  - RVG
  - AppDiskGroupName
  - DiskGroupName
- Sets the FireDrill attribute of the following resources to true:
  - IP
  - Lanman
  - RegRep
- Sets the ForFireDrill attribute of the following resources to `true` in the fire drill service group:
  - MountV
  - VMDg

This indicates that the volume being monitored by the VVRSnap agent belongs to the fire drill disk group.

About running the fire drill

The Fire Drill Wizard brings the fire drill service group online. Optionally, you can also run the fire drill using the Veritas InfoScale Operations Manager console.

In a Volume Replicator environment, when running the fire drill, the VVRSnap agent 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

About restoring the fire drill configuration

The Fire Drill Wizard takes the fire drill service group offline. Optionally, you can also restore the fire drill using the Veritas InfoScale Operations Manager console.

In a Volume Replicator environment, restoring the fire drill system to a prepared state, the VVRSnap agent does the following:

- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

About deleting the fire drill configuration

In a Volume Replicator environment, when deleting the fire drill configuration, the wizard does the following:

- Sets the FireDrill attribute of the following resources to false:
  - IP
  - Lanman
  - RegRep
- Unlinks the fire drill service group
- 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 Fire Drill Wizard performs the following basic operations in all replication environments:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
- Runs the fire drill by bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration by taking the fire drill service group offline
- Deletes the fire drill service group and any associated registry entries

See "About the Fire Drill Wizard" on page 424.

In Hitachi TrueCopy or EMC SRDF replication environments, the Fire Drill 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.

About 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. It links the fire drill service group to the corresponding application service group.
- In an HTC or SRDF environment, the wizard configures the Snap resource and sets the following attributes to the value `1`, which indicates:
  - **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.
- In an EMC SRDF environment, the wizard sets the following attribute values:
  - It sets **CopyMode** to one of the following, which indicates:
    - **Mirror**: Use the TimeFinder Mirror technology to create snapshots.
    - **Clone**: Use the TimeFinder Clone technology to create snapshots.
    - **Snap**: Use the TimeFinder Snap technology to create snapshots.
  - When the TimeFinder Clone technology is used, it sets **UseTgt** to one of the following, which indicates:

  - **0**: Use BCV devices to create snapshots.
  - **1**: Use STD devices to create snapshots.
- When the TimeFinder Snap technology is used, if a custom save pool area name is specified, it sets **SavePoolName** accordingly. The specified save pool area is used to create snapshots.
  If no value is specified on the SRDFSnap Resource Configuration panel, the default save pool area is used.

For information about the actual procedure:

See "Preparing the fire drill configuration" on page 436.

About 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 snapshot or 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, STD, or VDEV device
- Resumes replication
- Modifies the disk group name in the snapshot

For information about the actual procedure:

See "Running a fire drill" on page 443.

About restoring the fire drill configuration

When restoring the fire drill configuration to a prepared state, the wizard takes the fire drill service group offline, thereby taking the SRDF and HTC Snap agents offline.

This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

For information about the actual procedure:

See "Restoring the fire drill system to a prepared state" on page 448.

About deleting the fire drill configuration

When deleting the fire drill configuration, the wizard does the following:

- Delinks the fire drill service group from the corresponding application service group.
- Deletes the fire drill service group
- Deletes any associated registry entries

If you want to remove the hardware mirrors, you must do so manually.

For information about the actual procedure:

See "Deleting the fire drill configuration" on page 449.

For more information about the Hitachi TrueCopy Snap agent functions, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *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

1. Create a `.bat` file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\
PowerShell.exe -command "$ScriptName"
```

   In this entry, `$ScriptName` is either the fully qualified `.ps1` script, or the cmdlet entered by user.

   For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\
PowerShell.exe -command C:\myTest.ps1
```

2. 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 re-create 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.

The following table gives more details of the process of configuring and running fire drills with the wizard.

Table 15-1 Tasks 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 433.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See "Preparing the fire drill configuration" on page 436.

Table 15-1 Tasks for configuring and running fire drills *(continued)*

Action	Description
Re-create 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 re-create the fire drill configuration before running the fire drill. See "Re-creating a fire drill configuration that has changed" on page 445.
Run the fire drill	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. See "Running a fire drill" on page 443. Perform your own tests of the application to confirm that it is operational. **Note:** After completing the fire drill testing, run the wizard again as soon as possible to restore the configuration. Otherwise the fire drill service groups remain online. It is recommended that you restore a fire drill service group to a prepared state before running a fire drill on another service group.
Restore the fire drill configuration to a prepared state	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. This is a required action after running the fire drill. See "Restoring the fire drill system to a prepared state" on page 448. This operation takes the fire drill service group offline and reattaches snapshot mirrors.

Table 15-1 Tasks for configuring and running fire drills *(continued)*

Action	Description
Delete the fire drill configuration	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. See "Deleting the fire drill configuration" on page 449. The wizard deletes the service group on the secondary site. In a Volume Replicator 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. 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.

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 an application service group using IPv4 addresses, for each IP address in the service group, an IP address must be available to use on the secondary site for the fire drill service group.

  To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed. 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. More information on editing service group resources is available.

  See the *Cluster Server Administrator's Guide*.

  For IPv6, the IP address will be autogenerated.

- 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 Volume Replicator environment" on page 434.

See "Prerequisites for a fire drill in a Hitachi TrueCopy environment" on page 435.

See "Prerequisites for a fire drill in an EMC SRDF environment" on page 435.

Prerequisites for a fire drill in a Volume Replicator 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 433.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Volume Replicator environment:

- The primary and secondary sites must be fully configured with Volume Replicator replication and the global cluster option.
- 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 Volume Replicator 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 433.

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. Make sure that you have configured 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 433.

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. Make sure that you have configured disaster recovery with EMC SRDF.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.

- To take snapshots of R2 devices, appropriate additional devices must be associated with the RDF2 device group and fully established with the devices.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. Depending on the snapshot technology in use, this process involves the following tasks:
  - Mirror: Associate Symmetric Business Continuance Volumes (BCVs) and synchronize them with the secondary site source (STD devices).
  - Clone: Make sure that no clone session is in progress.
    The source and target devices must be of the exact same size.
  - Snap: Make sure that sufficient save pool area is configured and that no snap session is in progress.
    The source and target devices must be of the exact same size.
- 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 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 BCV, STD, or VDEV devices 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 Volume Replicator 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 433.

To prepare the fire drill configuration

1. Open the Solutions Configuration Center (From **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu).
2. In the Welcome panel, review the information and click **Next**.
3. In the System Selection panel, specify a system in the primary site cluster and click **Next**.

   See "System Selection panel details" on page 438.
4. 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 439.
5. 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 439.
6. 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 433.

   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 443.
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 448.
If the Re-create 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 re-create the fire drill configuration to bring it up to date. See "Re-creating a fire drill configuration that has changed" on page 445. Or you can clear the check box to re-create the configuration and run the fire drill on the existing configuration.

7 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 439.

8 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.

Volume Replicator replication	Disk Selection panel See "Disk Selection panel details" on page 439.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See "Hitachi TrueCopy Path Information panel details" on page 440. HTCSnap Resource Configuration panel See "HTCSnap Resource Configuration panel details" on page 441.
EMC SRDF replication	SRDFSnap Resource Configuration panel See "SRDFSnap Resource Configuration panel details" on page 441.

Click **Next**.

9 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.

See "Fire Drill Preparation panel details" on page 442.

When preparation is complete, click Next.

10 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 443.

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.

See "Preparing the fire drill configuration" on page 436.

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.

See "Preparing the fire drill configuration" on page 436.

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.

See "Preparing the fire drill configuration" on page 436.

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.

For IPv4, you must manually assign the IP address. For IPv6, the IP address will be autogenerated and displayed in the Virtual IP address field.

If the service group contains more than one IP and Lanman resource, this panel does not appear. 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.

See "Preparing the fire drill configuration" on page 436.

Disk Selection panel details

During fire drill preparation in a Volume Replicator 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 re-creating 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 FDnn.
Disk	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. If the production volumes reside on disks in the same disk group, you can store multiple snapshot volumes on a single disk. If the volumes in a disk group are configured on multiple RVG resources, provide a separate disk for each RVG. **Note:** The Fire Drill Wizard does not allow creating mirrors of multiple RVGs from a single disk group on the same disk. You must select a different disk for each RVG in a disk group. 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.
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This is a display-only field.

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.

See "Preparing the fire drill configuration" on page 436.

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 428.

SRDFSnap Resource Configuration panel details

Depending on the snapshot technology in use, the wizard validates the following when preparing for fire drill in an EMC SRDF replication environment:

- Mirror:
  The number of BCV devices matches that of the STD devices.
  The BCV devices are associated and synchronized with the STD devices.
- Clone:
  The number of BCV devices (or STD devices in case of Targets) matches that of the STD devices.
  No clone session is in progress.
- Snap:

The number of VDEV devices matches that of the STD devices.

No snap session is in progress.

If these criteria are not satisfied, the wizard displays a warning on this panel. The wizard does not check whether the sizes of the source and target devices match, and therefore does not display a warning. The following figure depicts such a warning.

However, you can proceed with the configuration. The wizards configures the fire drill service group, but is unable to bring the service group online.

This panel lists all the SRDFSnap resources that will be configured. If you do not want to include the dependent disk groups of a SRDFSnap resource in the fire drill, clear the check box against its name.

The name of the resource that is managing the LUNs that you want to snapshot appears as the Target Resource Name. 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.

You can specify the TimeFinder snapshot technology to be used for configuring fire drill for the SRDFSnap resources:

- **Mirror**
  BCV devices are used to create snapshots.
- **Clone**
  BCV devices are used to create snapshots. Optionally, you can specify that Target devices be used. If you select the **Use Target Devices** check box, STD devices are used to create snapshots.
- **Snap**
  VDEV devices are used to create snapshots. The default SavePoolArea is used. Optionally, to use a different SavePoolArea, specify its name.

To discover the most recent SRDF configuration information, click **Refresh**.

See "Preparing the fire drill configuration" on page 436.

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 428.

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 Volume Replicator 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 re-creating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See "Re-creating a fire drill configuration that has changed" on page 445.

See "Preparing the fire drill configuration" on page 436.

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 430.

For details on the operations that occur when running a fire drill, refer to the following topics:

- See "About Fire Drill Wizard operations in a Volume Replicator environment" on page 425.
- See "About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment" on page 428.

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 448.

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 re-creating a fire drill service group, go to step 6.

   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 Re-create Fire Drill Service Group panel, showing the differences.

   Choose one of the following:

   - Leave the option checked to re-create the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.
     See "Re-creating a fire drill configuration that has changed" on page 445.
   - To run the fire drill on the existing configuration, clear the option to re-create 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 430.

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 448.

Re-creating 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. The wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Re-create Fire Drill Service Group panel.

The wizard also checks the RVGs configured for disk groups. If a single RVG is configured per disk, the wizard allows you to re-create the service group; the existing snapshots are retained. If multiple RVGs are configured on a disk, the wizard only allows you to delete the service group; the existing snapshots are deleted. To create a corresponding new one, you need to launch the wizard again and perform the fire drill preparation steps.

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 re-create 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.

You have the following choices from the Re-create Fire Drill Service Group panel:

- Leave the option checked to re-create the fire drill service group. Proceed with using the wizard to re-create the configuration to match the application service group. The wizard deletes the existing fire drill configuration first, before creating the new one.

  For a Volume Replicator replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it re-creates 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 re-create 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.

The following procedure describes the choice of re-creating the fire drill configuration.

Note: Veritas recommends that you do not use this procedure to re-create any existing fire drill service groups after performing an upgrade. Instead, use the Fire Drill Wizard to delete the existing service groups and create corresponding new ones.

To re-create the fire drill configuration if the service group has changed

1. In the Re-create Fire Drill Service Group panel, leave the option checked to re-create the configuration before running the fire drill.

   For a Volume Replicator 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 Volume Replicator 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 433.

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.

Volume Replicator 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 439.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See "Hitachi TrueCopy Path Information panel details" on page 440. HTCSnap Resource Configuration panel See "HTCSnap Resource Configuration panel details" on page 441.
EMC SRDF replication	SRDFSnap Resource Configuration panel See "SRDFSnap Resource Configuration panel details" on page 441.

Click **Next**.

6 The Fire Drill Preparation panel is displayed. Wait while the wizard re-creates the fire drill service group.

For Volume Replicator 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.

7 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 443.

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:

- See "About Fire Drill Wizard operations in a Volume Replicator environment" on page 425.
- See "About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment" on page 428.

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 6.

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 Volume Replicator 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 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. If the wizard detects that the fire drill service group is different from the application service group, it displays the Re-create Fire Drill Service Group panel. Clear the option to re-create 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**.

Fire Drill Deletion panel details

The Fire Drill Deletion panel shows the progress of the deletion tasks. Wait until all tasks are complete before you click **Next** to proceed to the Summary panel and exit the wizard.

See "Deleting the fire drill configuration" on page 449.

The Fire Drill Deletion panel also appears if you selected the option to re-create a fire drill configuration. During a re-create operation, wait until all deletion tasks are complete and then click **Next** to continue with the fire drill preparation.

When the wizard re-creates a configuration in a Volume Replicator environment, it deletes the service group but maintains the required snapshot volumes. If you close the wizard without continuing to the preparation step, the snapshot volume information is lost and the fire drill configuration is fully deleted.

Considerations for switching over fire drill service groups

In a Volume Replicator environment, if you directly switch the fire drill service group from one node to another, the VVRSnap resource fails to come online on the target node. The fire drill service group depends on the RVG service group. To make the switch successfully, you must first switch the RVG service group to the intended node and then switch the fire drill service group.

Section 7

Microsoft Clustering Solutions

- Chapter 16. Microsoft clustering solutions overview
- Chapter 17. Deploying SFW with Microsoft failover clustering
- Chapter 18. Deploying SFW with Microsoft failover clustering in a campus cluster
- Chapter 19. Deploying SFW and VVR with Microsoft failover clustering

Chapter 16

Microsoft clustering solutions overview

This chapter includes the following topics:

- About Microsoft clustering with high availability
- About Microsoft clustering with Volume Replicator
- About Microsoft clustering with campus clustering
- About the SFW-Microsoft clustering-Volume Replicator configuration

About Microsoft clustering with high availability

Microsoft clustering may be used with Storage Foundation 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 Volume Replicator

Microsoft clustering may be used with Storage Foundation and Volume Replicator to provide replication support for your application. Using Volume Replicator 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 Volume Replicator 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 Veritas InfoScale 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-Volume Replicator 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 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-Volume Replicator 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-Volume Replicator configuration with Microsoft clustering. For a SFW-Microsoft clustering-Volume Replicator 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-Volume Replicator configuration

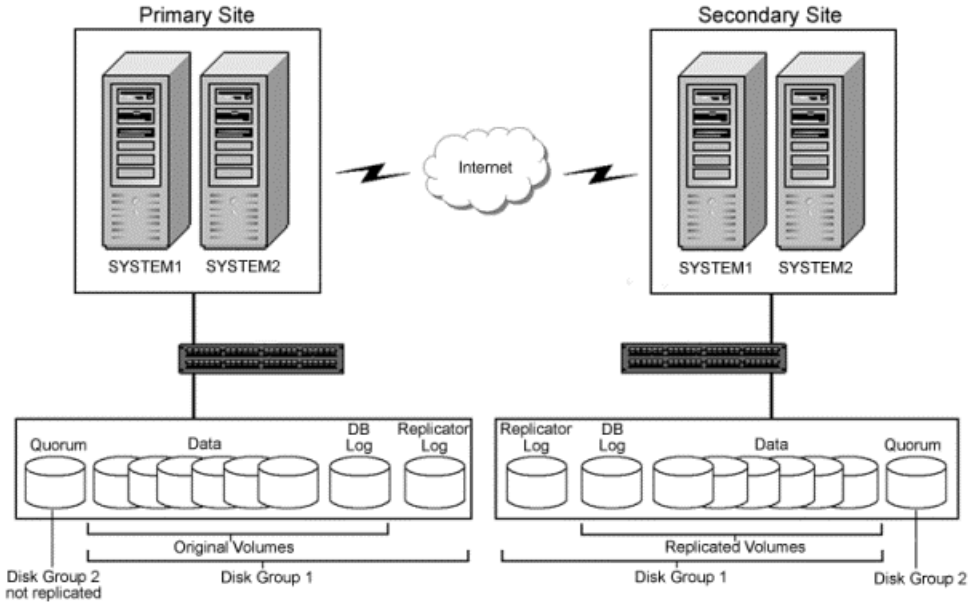

The quorum volume is not replicated from the primary site to the secondary site. Each site has its own quorum volume. A two-way or four-way mirror is recommended for the quorum volume for redundancy.

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.

Chapter 17

Deploying SFW with Microsoft failover clustering

This chapter includes the following topics:

- Tasks for deploying InfoScale Storage with Microsoft failover clustering
- Reviewing the configuration
- Configuring the storage hardware and network
- Establishing a Microsoft failover cluster
- Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering
- Creating SFW disk groups and volumes
- Creating a group for the application in the failover cluster
- Installing the application on cluster nodes
- Completing the setup of the application group in the failover cluster
- Implementing a dynamic quorum resource
- Verifying the cluster configuration
- Configuring InfoScale Storage in an existing Microsoft Failover Cluster

Tasks for deploying InfoScale Storage with Microsoft failover clustering

This chapter describes how to install and configure InfoScale Storage with Microsoft failover clustering in a new installation, using a two-node active/passive cluster configuration as an example. 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:

- *Storage Foundation and High Availability Solutions Microsoft Clustering Solutions Guide for Microsoft Exchange*
- *Storage Foundation and High Availability Solutions Microsoft Clustering Solutions Guide for Microsoft SQL*

The following table outlines the high-level objectives for implementing the configuration and the tasks for each objective:

Table 17-1 Task list for deploying InfoScale Storage with Microsoft failover clustering

Objectives	Tasks
See "Reviewing the configuration" on page 459.	Understanding the configuration for the failover cluster
See "Configuring the storage hardware and network" on page 498.	■ Install the operating system on both nodes. ■ Make necessary networking settings on both nodes.
See "Establishing a Microsoft failover cluster" on page 529.	Refer to Microsoft documentation for instructions on establishing a cluster under Microsoft failover clustering.
See "Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering" on page 502.	■ Install InfoScale Storage. ■ Install InfoScale Storage Option for Microsoft Cluster Service (MSCS).
See "Creating SFW disk groups and volumes" on page 472.	■ 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 InfoScale Storage with Microsoft failover clustering *(continued)*

Objectives	Tasks
See "Creating a group for the application in the failover cluster" on page 478.	■ 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.
See "Installing the application on cluster nodes" on page 479.	■ 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.
See "Completing the setup of the application group in the failover cluster" on page 480.	■ 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.
See "Implementing a dynamic quorum resource" on page 481.	■ 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.
See "Verifying the cluster configuration" on page 520.	■ 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 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 Microsoft failover clustering and two servers

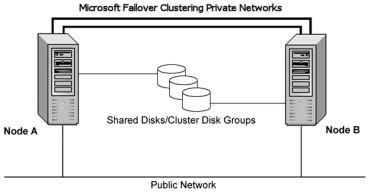

This configuration does not include DMP. For information about DMP and clustering:

See "Overview of configuration tasks for adding DMP DSMs" on page 178.

Key points about the configuration:

- A Microsoft failover cluster must be running before you install InfoScale Storage. Therefore, you need to set up the hardware and install the operating system on both systems and establish the failover cluster before installing InfoScale Storage. Installing InfoScale Storage requires a reboot, but a reboot on the active cluster node causes it to fail over. Use a "rolling install" procedure to install InfoScale Storage 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, Veritas 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 these procedures 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, Veritas 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. Verify that each system can access the storage devices.
4. 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 by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.
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 that the public network adapter is the first bound adapter by following these steps sequentially:
   - In the Network Connections window, click **Advanced > Advanced Settings**.
   - In the Adapters and Bindings tab, verify that 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 On the Public Properties dialog box, on the General tab:

- Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
- Click **Properties**.

8 Select the **Use the following DNS server addresses** option.

9 Verify the correct value for the IP address of the DNS server.

10 Click **Advanced**.

11 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.

12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.

13 Close the window.

Establishing a Microsoft failover cluster

Before installing InfoScale Storage, you must first verify that Microsoft failover clustering is enabled (on a new Windows Server installation), 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**).

   On Windows 2012 operating systems, launch **Administrative Tools** from the **Start** screen.
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. Veritas 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 On 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.

Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering

This section assumes you are running a Microsoft failover cluster and you are installing InfoScale Storage on an inactive system that does not own any cluster resources.

Veritas recommends that you perform a rolling installation of InfoScale Foundation or InfoScale Storage. For a rolling installation, you must first install the product on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After the product is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then, install the product on the other inactive systems in the Microsoft failover cluster simultaneously.

Perform the tasks that are described in the following topics:

- See "Pre-installation task: moving the online groups" on page 502.
- See "About installing the Veritas InfoScale products" on page 503.
- See "Post-installation task: moving the online groups" on page 503.

Pre-installation task: moving the online groups

If your resource groups are on the system where you are installing InfoScale Storage, you must move the resource groups from the system to another system in the cluster.

To move the online groups

1. Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).

   On Windows 2012 operating systems, use the **Start** screen.

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 *nodeName***.

   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.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Installing the server components using the installation wizard

The product installation wizard allows you to install the product on multiple systems at a time.

Before you begin to install the product ensure that you have reviewed the installation prerequisites, licensing, and the product co-existence details.

Note: If you plan to install InfoScale Storage in an active Microsoft Failover Cluster, ensure that you have reviewed the applicable pre-requisites, and use the "rolling-install" procedure to perform the product installation. To use the "rolling-install" procedure, install InfoScale Storage first on the inactive cluster node. Then move the cluster resources to the other node and install the product on the now inactive node.

Perform the following steps to install the server components

1. Download the installation package from the following location:

   https://sort.veritas.com

2. Allow the autorun feature to start the installation or double-click **Setup.exe**.

   The CD browser appears.

   Note: If you install the software using the product software disc, the CD browser displays the installation options for all the products that are specified earlier. However, if you download the installation package from the Veritas website, the CD browser displays the installation options only for the product to be installed.

3. Click the required product-specific tab and then click the link to install the components.

   Note: The client components are installed by default along with the server components. However, the client components are not installed if the system is a server core machine.

   In addition to the product-specific tabs, the CD browser also provides the following links:

Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
Windows Data Collector	Click to verify that your configuration meets all pertinent software and hardware requirements.
SORT	Click to access the Veritas Services and Operations Readiness Tools (SORT) site. In addition to the product download you can also download the custom reports about your computer and Veritas enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
Browse Contents	Click to view the software disc contents.
Technical Support	Click to contact Veritas Technical Support.

4 On the Welcome panel, review the list of prerequisites and click **Next**.

Note that the **Check for product updates** check box is selected by default. The wizard searches for the available product updates on the SORT website. You can then download and apply the available updates. If you do not want to apply the available updates, clear the selection of **Check for product updates** check box.

5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.

The **Participate in the Veritas Product Improvement Program by submitting system and usage information anonymously** check box is selected by default. The Product Improvement Program allows the wizard to collect installation, deployment, and usage data and submit it anonymously to Veritas. The collected information helps identify how customers deploy and use the product. If you do not want to participate in the Product Improvement Program, clear the selection of the check box.

6 On the Product Updates panel, review the list of available product updates.

This panel appears only if you have selected the **Check for product updates** check box on the Welcome panel.

The product updates comprise of the pre-installation patches, post-installation patches, high availability Agents, and array-specific modules. The panel lists the available pre-installation patches and the post-installation patches. Download and apply the pre-installation patches in the sequence shown in the table and rerun the wizard. After the successful installation of the product, apply the post-installation patches. Also download and install the high availability agents and array-specific modules from the SORT website.

7 On the System Selection panel, select the systems and the desired Installation and Product options:

You can select the systems in one of the following ways:

- In the System Name or IP text box, manually type the system name or its IP address and click **Add**.

  Note: The wizard does not support the Internet Protocol version 6. To add the systems having Internet protocol version 6, you must type the system name.

  The local host is populated by default.

- Alternatively, browse to select the systems.

The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks, and notes the details in the Verification Details box. To review the details, select the desired system.

To select the installation and product options, perform the following tasks on each of the selected system.

Note: To apply the selection to multiple systems, select the system for which you have selected the installation and product options and then click **Apply to multiple systems**.

See "Applying the selected installation and product options to multiple systems" on page 470.

- By default the wizard uses %ProgramFiles%\Veritas as the installation directory. To customize the installation directory, click **Browse** and select the desired location. Click **OK**.

  Install the product at the same location on all the systems.

  If you are upgrading the product, the installation directory is selected by default.

  Note: The installation directory must contain only English characters, if:

  - You plan to configure the cluster for single sign-on authentication.

  - Your system runs a non-English locale operating system.

- Select the required license type from the **License key** drop-down list.

  Note: The default license type is "Keyless".

  If you select the "Keyless" license type, all the available product options are displayed and are selected by default.

  If you select "User entered license key" as your license type, the License Details panel appears by default. On the License Details panel, enter the license key and then click **Add**. You can add multiple licenses for the various product options you want to use.

  The wizard validates the entered license keys and displays the relevant error if the validation fails. After the validation is complete, click **OK**.

- From the list of product options, select the options to be installed.
  The options differ depending on the product you install.

8 On the System Selection panel, click **Next**.

Note that the wizard fails to proceed with the installation, unless all the selected systems have passed the validation checks and are ready for installation. In case the validation checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the installation, select the system and click **Re-verify** to re-initiate the validation checks for this system.

9 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This selection reboots all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

10 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

If an installation is not successful on any of the systems, the status screen shows a failed installation.

Note: During the upgrade, the Installation panel displays a list of services and processes running on the systems. Select a system to view the services and processes running on it and review the list.

The wizard stops the product-specific services and discovers the processes running, if any, on the systems. These processes need to be stopped to proceed with the operation. Click **Next** to forcefully stop the processes and proceed with the operation. Alternatively, you can manually stop the processes. If the services or processes cannot be stopped, the operation fails. Rectify the error and then click **Retry** to validate the affected system again. Click **Retry All** to validate all the systems again.

In case you want to proceed with the upgrade without stopping a particular process, contact Veritas Technical Support.

---

11 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed on any of the system, refer to the log file for details. You may have to re-install the software.

12 On the Finish panel, click **Finish**.

If you had chosen to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

In case you had not selected to initiate the auto reboot, ensure that you manually reboot these systems.

This completes the product installation. You must now proceed to configure the required components. Refer to the component-specific guides for more details about the configuration tasks.

Note: If you have installed InfoScale Storage with Microsoft Failover Cluster, but if Microsoft failover cluster is not yet configured, you must register the InfoScale Storage resources, after configuring the Microsoft failover cluster software.

See "Registering the InfoScale Storage resource DLLs" on page 471.

However, if you have installed InfoScale Storage in an active Microsoft Failover Cluster, then you must remove the physical disk resources for all the basic disks. You must do so before configuring the SFW cluster disk groups. Failing this, a reservation conflict occurs.

Applying the selected installation and product options to multiple systems

To apply the selected installation and product options to multiple systems, perform the following steps:

1 Click on any one of the selected systems and select the desired installation and product options.

2 Click **Apply to multiple systems**.

3 On the Apply Installation Options panel, select the installation options to be applied and then select the desired systems. Click **OK**.

Note: The installation directory is selected by default on the systems where the product is being upgraded. The selected **Install Directory** option does not apply to these systems.

Registering the InfoScale Storage resource DLLs

You must perform this task only if you have installed InfoScale Storage with Microsoft failover cluster option, but Microsoft failover cluster is not yet configured in your environment.

- Using Windows Powershell cmdlets:
  - Import the FailoverClusters module
    Type the following cmdlet:
    ```
    Import-module failoverclusters
    ```
  - Register the Volume Manager Disk Group (VMDg) resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Volume Manager Disk Group"
    C:\Windows\Cluster\vxres.dll -DisplayName "Volume Manager Disk
    Group"
    ```
  - Register the Replicated Volume Group (RVG) resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Replicated Volume Group"
    C:\Windows\Cluster\mscsrvgresource.dll -DisplayName "Replicated
    Volume Group"
    ```
  - Register the Volume Manager Shared Volume resource type
    Type the following cmdlet:
    ```
    Add-ClusterResourceType "Volume Manager Shared Volume"
    C:\Windows\Cluster\vxvolres.dll -DisplayName "Volume Manager
    Shared Volume"
    ```

Installing the client components

Note: Client components cannot be installed on server core systems.

Perform the following steps to install the client components

1. Log on to any node in the cluster.
2. From the Veritas Services and Operations Readiness Tools (SORT) site, download the client-specific .msi file.

   https://sort.veritas.com/

msi name	Client component name
vrtsvcscscm.msi	VCS Cluster Manager (Java Console)
VRTSobgui.msi	Veritas Enterprise Administrator (VEA)

3. Follow the instructions that the install program provides, to complete the installation.

Post-installation task: moving the online groups

You can move the resource groups from the current system, back to the previous system where InfoScale Storage is installed.

To move the online groups

1. Open the Failover Cluster Management tool (**Start > Administrative Tools > Failover Cluster Management**).

   On Windows 2012 operating systems, use the **Start** screen.
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 *nodeName***.

   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.

Creating SFW disk groups and volumes

Use Storage Foundation 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:

- See "Planning disk groups and volumes" on page 473.
- See "Creating a dynamic disk group" on page 508.
- See "Creating dynamic volumes" on page 475.

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 SQL Server, 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 481.
  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 Microsoft failover clustered database with disks for data, the log, and the quorum resource

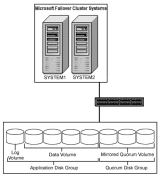

Creating a dynamic disk group

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.

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

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

To create a dynamic disk group

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:

- Enter the disk group name (for example, **DG1**).
- Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
- 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.
- Click **Next**.

Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

---

7 Click **Next** to accept the confirmation screen with the selected disks.

8 Click **Finish** to create the dynamic 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.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

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

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. For Site Preference, leave the setting as **Siteless** (the default).
- Automatic disk selection is the default setting. To manually select the disks, click **Manually select disks** 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.
- Click **Next**.

7 Specify the volume attributes.

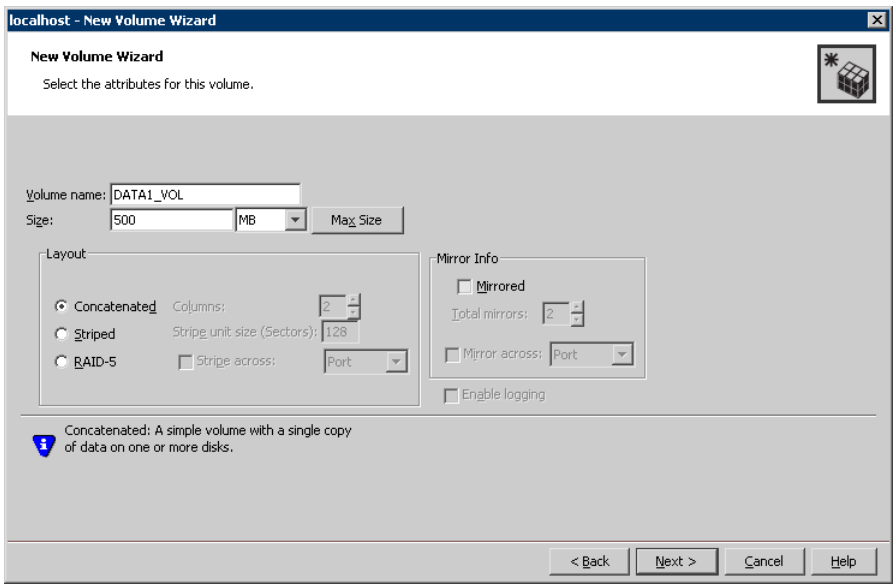

- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.

- 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.
- Click **Next**.

8 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 disk.
- Click **Next**.

9 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**.

10 Click **Finish** to create the new volume.

11 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 480.

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.

   On Windows 2012 operating systems, use the **Start** screen.

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 the next step.

   To launch Failover Cluster Management, select **Start > Administrative Tools > Failover Cluster Management**.

   On Windows 2012 operating systems, use the **Start** screen.

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 earlier:
  See "Creating dynamic volumes" on page 475.

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 17-3 Application group dependencies

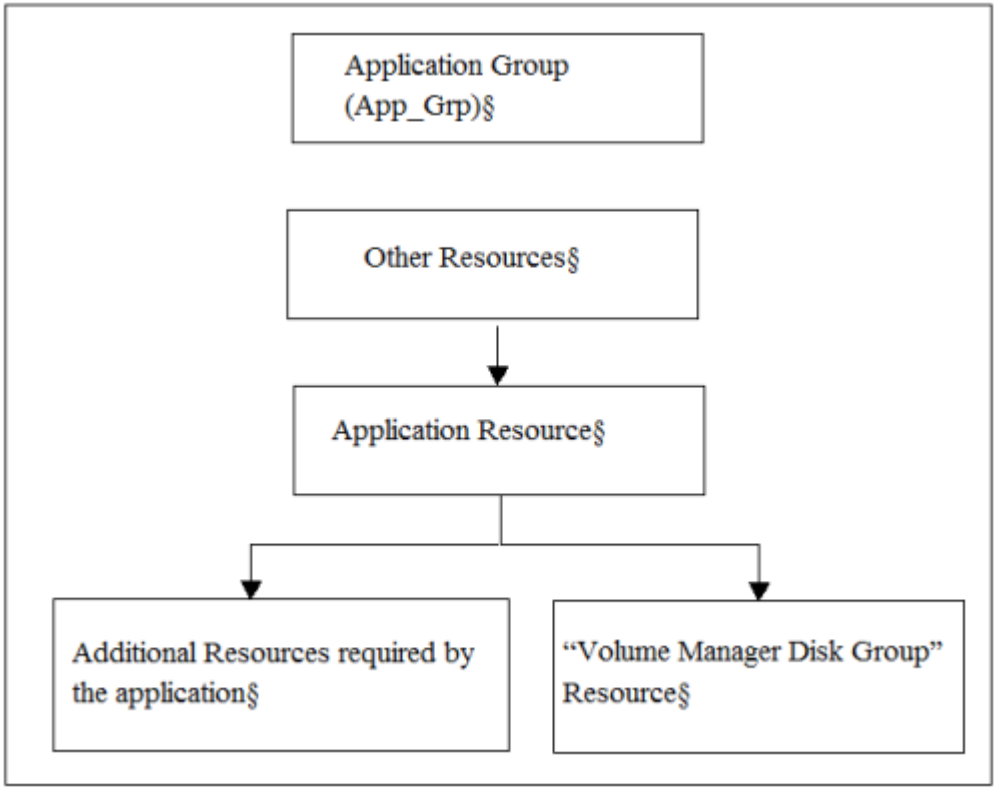

- Testing: After the application group is set up, test it by moving the cluster resources to another node and then move them back.

Implementing a dynamic quorum resource

Although Veritas 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:

- See "Creating a dynamic cluster disk group and a mirrored volume for the quorum resource" on page 482.
- See "Adding a VMDg resource for the quorum" on page 515.
- See "Changing the quorum resource to a dynamic mirrored quorum resource" on page 516.

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. Veritas 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 472.

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 a VMDg resource for the quorum

You 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

1. If Failover Cluster Management is already open, then proceed to the next step.

   To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.

   On Windows 2012 operating systems, use the **Start** screen.

2. Verify that the cluster is online on the same node where you created the disk group.

3. In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
4. Right-click the new group and rename it, for example **QUORUM**.
5. Right-click **QUORUM** and select **Add a resource > More resources > Add Volume Manager Disk Group**.
6. Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
7. In the General tab of the Properties dialog box, type a name for the resource in the **Resource Name** field, for example, **QUORUM_DG_RES**.
8. On the Properties tab, in the **Disk Group Name** field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
9. Right-click the Quorum disk group resource (for example, QUORUM_DG_RES) in the left pane and select **Bring this resource online**.

   The specified disk group resource, QUORUM_DG_RES resource, is created under the Quorum group (for example, QUORUM).

Changing the quorum resource to a dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

1. In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.
2. The Configure Cluster Quorum Wizard opens. Review the screen and click **Next**.
3. Select either the **Node and Disk Majority** or No **Majority: Disk Only** radio button, and click **Next**.
4. Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.

   This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, **QUORUM_DG_RES**.
5. Review the information in the Confirmation screen and click **Next**.
6. Click **Finish** to close the wizard.

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, as follows:

- 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.

Caution: Do not simulate failover in a production environment.

To move online groups

1. Open the Failover Cluster Management tool by clicking **Start > Administrative Tools > Failover Cluster Management**.

   On Windows 2012 operating systems, use the **Start** screen.

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 *nameOfNode***.

   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 > Administrative Tools > Failover Cluster Management** from any node in the cluster.

   On Windows 2012 operating systems, use the **Start** screen.

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 the first step, select the resource group, and use **Move this service or application to another node > Move to node *nameOfNode*** to move the resource group.

Configuring InfoScale Storage in an existing Microsoft Failover Cluster

After you have configured an application for high availability, in a Microsoft Failover Cluster, you may want to move the application data from the existing disks to the InfoScale Storage -controlled storage disks. This task involves installing InfoScale Storage on all the cluster systems, converting the already configured basic disks to dynamic disks, and then adding the dynamic disk group resource to the application role.

Notes:

- You are required to reboot the systems to successfully install InfoScale Storage. Also, you need take the application role offline before you begin to convert the basic disks to dynamic disks. These procedures result in application down-time.
- For the steps performed using the Failover Cluster Manager, refer to the Microsoft documentation for details. For the steps performed using VEA, refer to the SFW administrator's guide for the details about Disk and Volume tasks.

To configure InfoScale Storage in an existing Microsoft Failover Cluster, perform the following steps:

1 Install InfoScale Storage on all the cluster systems. You must select the Microsoft Failover Cluster option during the installation.

Note: At the end of the installation process you are required to reboot the system. To ensure less down-time, you can first install InfoScale Storage on the cluster systems other than the one where the application role is online. After you complete the installation on all these systems initiate the installation on the system where the application role is online.

---

2 After the cluster systems have restarted, using the Failover Cluster Manager, stop the application role to bring the resources offline.

3 Check the resource dependencies to note the storage resource dependencies.

4. Remove the existing basic disk storage resource from the application role.
5. Remove the basic disk from the available storage.

   This step removes the basic disk resource from the application role and takes the disk offline.
6. Using VEA, bring the basic disk online.
7. Create a new clustered dynamic disk group using the basic disk that is brought online in step 6. Before you create a dynamic disk group, ensure that minimum 16MB free space is available in the disk. This space is required to upgrade a basic disk to a dynamic disk.
8. Using Failover Cluster Manager, to the application role, add a Volume Manager Disk Group resource for the dynamic disk group that is created in step 7.
9. Set the resource dependencies as noted in step 3.

   All the earlier storage resource dependencies must now be replaced with the Volume Manager Disk Group resource.
10. Bring the application role online.

Chapter 18

Deploying SFW with Microsoft failover clustering in a campus cluster

This chapter includes the following topics:

- Tasks for deploying InfoScale Storage with Microsoft failover clustering in a campus cluster
- Reviewing the configuration
- Configuring the storage hardware and network
- Establishing a Microsoft failover cluster
- Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering
- Creating disk groups and volumes
- Implementing a dynamic quorum resource
- Setting up a group for the application in the failover cluster
- Installing the application on the cluster nodes
- Completing the setup of the application group in the cluster
- Verifying the cluster configuration

Tasks for deploying InfoScale Storage with Microsoft failover clustering in a campus cluster

This chapter presents a Microsoft failover clustering example with a two-node campus cluster.

The table below outlines the high-level objectives and the tasks for each objective:

Table 18-1 Task list for deploying InfoScale Storage with Microsoft failover clustering in a campus cluster

Objectives	Tasks
See "Reviewing the configuration" on page 489.	■ Review the configuration requirements. ■ Overview of a campus cluster using Microsoft clustering and recovery scenarios.
See "Configuring the storage hardware and network" on page 498.	■ Install and configure the hardware for each node in the cluster. ■ Verify the DNS settings and binding order for all systems.
See "Establishing a Microsoft failover cluster" on page 529.	■ 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.
See "Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering" on page 502.	■ Install InfoScale Storage on Node A (Node B active). ■ Install InfoScale Storage on Node B (Node A active).
See "Creating disk groups and volumes" on page 504.	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.
See "Implementing a dynamic quorum resource" on page 514.	■ 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.

Table 18-1 Task list for deploying InfoScale Storage with Microsoft failover clustering in a campus cluster *(continued)*

Objectives	Tasks
See "Setting up a group for the application in the failover cluster" on page 517.	■ 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.
See "Installing the application on the cluster nodes" on page 518.	■ 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.
See "Completing the setup of the application group in the cluster" on page 519.	■ 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.
See "Verifying the cluster configuration" on page 520.	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 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, refer to the following topics:

- See "Overview of campus clustering with Microsoft clustering" on page 491.
- See "Campus cluster failure with Microsoft clustering scenarios" on page 492.

Figure 18-1 Campus clustering with Microsoft clustering configuration example

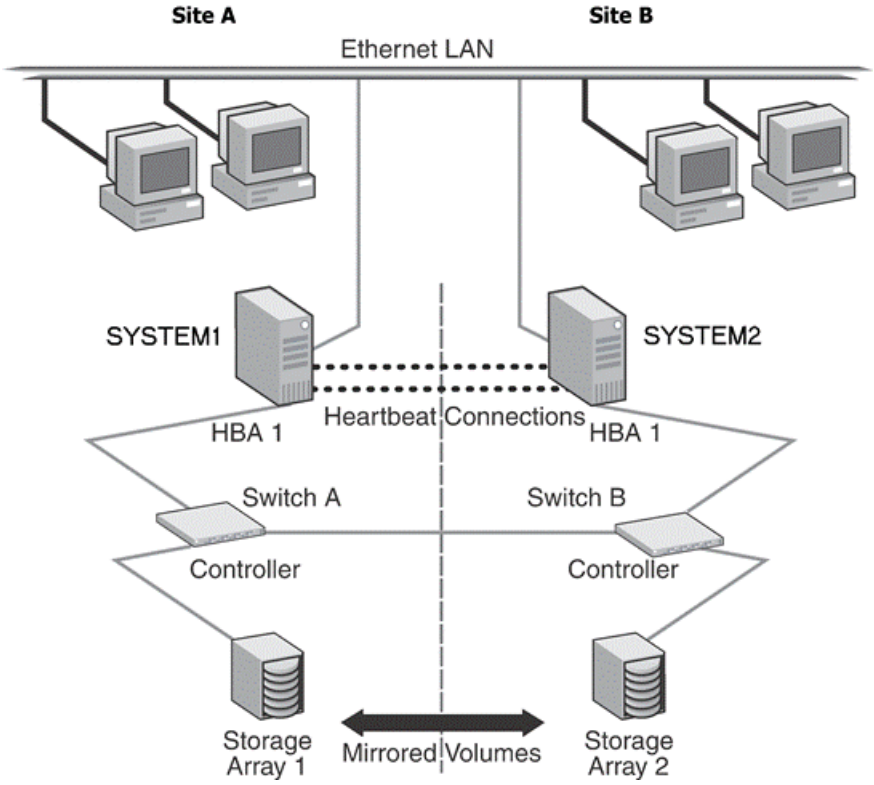

The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array and contains mirrored data of the storage on the other array. Each disk group should contain the same number of disks on each site for the mirrored volumes.

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 "Overview of configuration tasks for adding DMP DSMs" on page 178.

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

The following figure 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, 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 18-2 Typical campus clustering configuration with Microsoft clustering

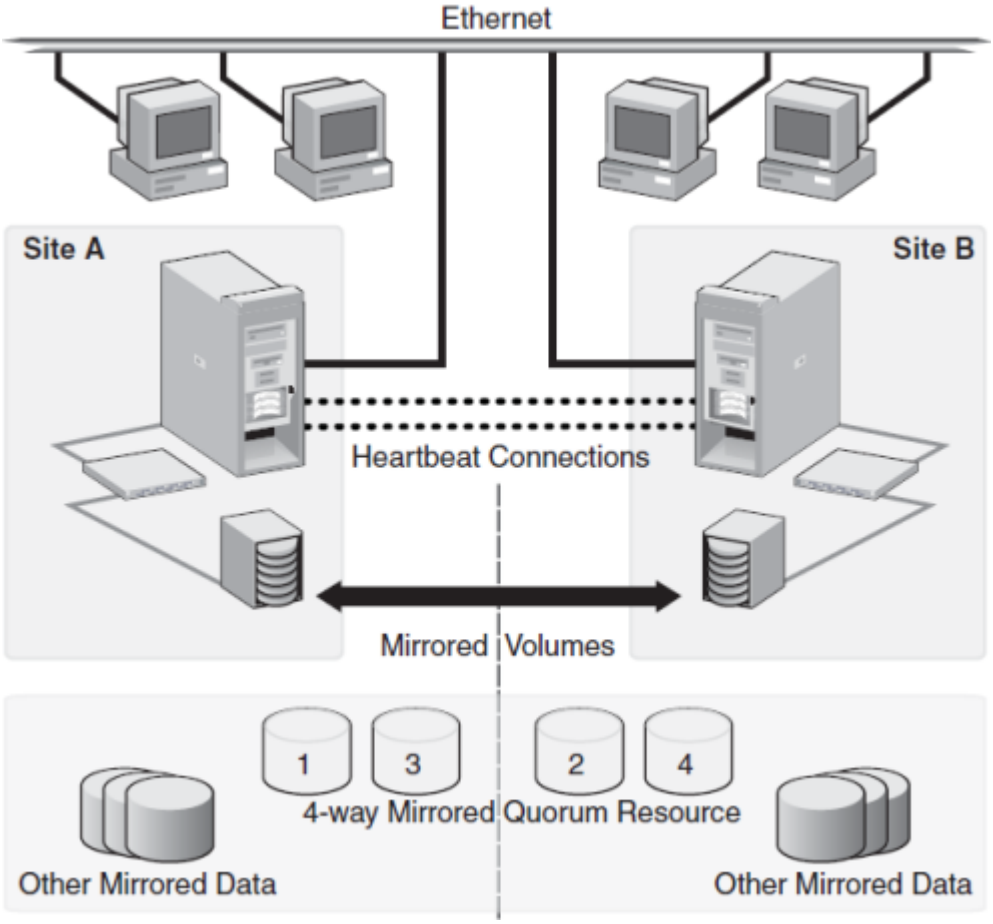

Most customers use hardware RAID to protect the quorum disk, but that will not work when a natural disaster takes down the primary node and its attached storage. If the quorum resource is lost to the cluster, the cluster fails, because none of the cluster servers can gain control of the quorum resource 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 496.

The following table lists failure situations and the outcomes that occur.

Table 18-2 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.

Table 18-2 List of failure situations and possible outcomes *(continued)*

Failure situation	Outcome	Comments
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.
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.

Table 18-2 List of failure situations and possible outcomes *(continued)*

Failure situation	Outcome	Comments
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.
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.

Table 18-2 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 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 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.

- See "Quorum" on page 497.
- See "Cluster ownership of the quorum resource" on page 497.
- See "The vxclus utility" on page 497.

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 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 provides the vxclus command line utility to allow forcing a failover to the secondary site. The `vxclus enable` command 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 storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat these procedures 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, Veritas 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. Verify that each system can access the storage devices.
4. 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 by clicking **Start > Control Panel**.

   On Windows 2012 operating systems, use the **Settings** menu from the **Start** screen.
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 that the public network adapter is the first bound adapter by following these steps sequentially:
   - In the Network Connections window, click **Advanced > Advanced Settings**.
   - In the Adapters and Bindings tab, verify that 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. On the Public Properties dialog box, on the General tab:
   - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
   - Click **Properties**.
8. Select the **Use the following DNS server addresses** option.
9. Verify the correct value for the IP address of the DNS server.

10 Click **Advanced**.

11 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.

12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

   To find the domain suffix, click **Start > Control Panel > System**. The domain suffix is listed in the "Computer Name, domain, and workgroup settings" section.

13 Close the window.

Establishing a Microsoft failover cluster

Before installing InfoScale Storage, you must first verify that Microsoft failover clustering is enabled (on a new Windows Server installation), 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**).

   On Windows 2012 operating systems, launch **Administrative Tools** from the **Start** screen.

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. Veritas 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. On the Confirmation screen, verify that the cluster configuration is correct, and then click **Next** to create the cluster.

11. Review the Summary page and then click **Finish** to close the wizard.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use these disks when you create your SFW cluster disk groups, you must first remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs. After creating the SFW cluster disk groups, you will add Volume Manager Disk Group resources to the cluster, instead of physical disk resources.

Connecting the two nodes

Make the necessary connections between the two sites. The cluster is already active on Server A, so Microsoft clustering is now in control of the cluster storage on Server A, and both nodes of the storage cannot be accessed at the same time by the operating system.

To connect the two nodes

1. Connect corresponding cables between the three network cards on the two sites.
2. Connect the two switches at the two sites through the storage interconnect.
3. Test the connectivity between the two sites. Test the IP addresses of all the network adapter cards in the cluster. Bring up the command window and type `ping ipaddress`, where the *ipaddress* is the corresponding network adapter in the other node.

Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering

This section assumes you are running a Microsoft failover cluster and you are installing InfoScale Storage on an inactive system that does not own any cluster resources.

Veritas recommends that you perform a rolling installation of InfoScale Foundation or InfoScale Storage. For a rolling installation, you must first install the product on an inactive system. Our example uses a two node configuration, so the inactive system is the second node. After the product is installed on an inactive system, move the resource groups to this system, and make the other systems inactive. Then, install the product on the other inactive systems in the Microsoft failover cluster simultaneously.

Perform the tasks that are described in the following topics:

- See "Pre-installation task: moving the online groups" on page 502.
- See "About installing the Veritas InfoScale products" on page 503.
- See "Post-installation task: moving the online groups" on page 503.

Pre-installation task: moving the online groups

If your resource groups are on the system where you are installing InfoScale Storage, you must move the resource groups from the system to another system in the cluster.

To move the online groups

1. Open Failover Cluster Management (**Start > Administrative Tools > Failover Cluster Management**).

   On Windows 2012 operating systems, use the **Start** screen.

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 *nodeName***.

   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.

About installing the Veritas InfoScale products

For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.

You can use Veritas InfoScale Operations Manager to monitor the status of the application. For more information, see the Veritas InfoScale Operations Manager product documentation.

Post-installation task: moving the online groups

You can move the resource groups from the current system, back to the previous system where InfoScale Storage is installed.

To move the online groups

1. Open the Failover Cluster Management tool (**Start > Administrative Tools > Failover Cluster Management**).

   On Windows 2012 operating systems, use the **Start** screen.

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 *nodeName***.

   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.

Creating disk groups and volumes

Use Storage Foundation 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.

Refer to the following topics:

- See "Example disk group and volume configuration in campus cluster" on page 505.
- See "Considerations when creating disks and volumes for campus clusters" on page 506.
- See "Viewing the available disk storage" on page 507.
- See "Creating a dynamic disk group" on page 508.
- See "Adding disks to campus cluster sites" on page 509.
- See "Creating volumes for campus clusters" on page 510.

Example disk group and volume configuration in campus cluster

The illustration that follows shows a typical Microsoft failover cluster with a campus cluster setup of disks. For campus clusters, each disk group must contain an equal number of disks on each site. 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 18-3 Microsoft failover cluster with campus cluster disks and disk groups example

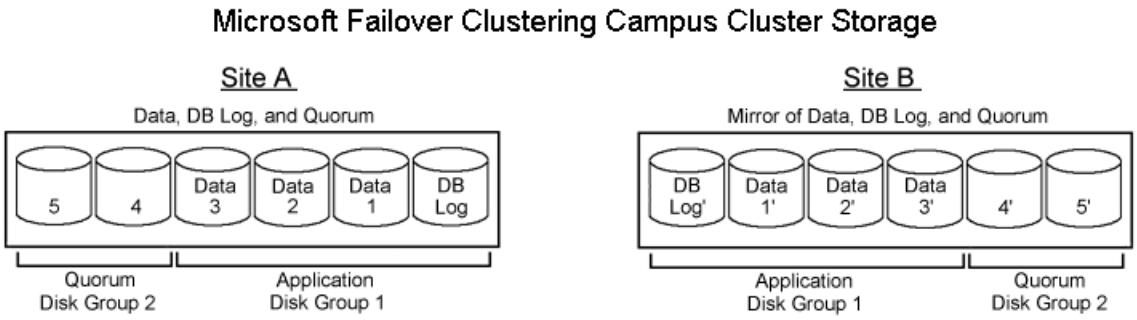

Considerations when creating disks and volumes for campus clusters

When you create the disk groups for a campus cluster, ensure that each disk group has the same number of disks on each physical site. You create each volume as 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.

Veritas recommends using the SFW site-aware allocation feature for campus cluster storage. Site-aware allocation can ensure that site boundary limits are maintained for operations like volume grow, subdisk move, and disk relocation.

Enabling site-aware allocation for campus clusters requires the following steps in the VEA:

- After creating the disk groups, you tag the disks with site names to enable site-aware allocation. This is a separate operation, referred to in the VEA as adding disks to a site.
  As an example, say you had a disk group with four disks. Disk1 and Disk2 are physically located on Site A. Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to "site_a" and add Disk3 and Disk4 to "site_b".
- During volume creation, you specify the volume site type as Site Separated. This ensures that the volume is restricted to the disks on the selected site.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

For more information on site-aware allocation, refer to the *Storage Foundation Administrator's Guide*.

When you create the volumes for a campus cluster, consider the following:

- During disk selection, configure the volume as "Site Separated" and select the two sites of the campus cluster from the site list.
- For volume attributes, select the "mirrored" and "mirrored across enclosures" options.
- Veritas recommends using either simple mirrored (concatenated) or striped mirrored options 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.
- During the volume creation procedure for Site Separated volumes, you can only create as many mirrors as there are sites. However, once volume creation is complete, you can add additional mirrors if desired.
- Choosing "Mirrored" and the "mirrored across" option without having two enclosures that meet requirements causes new volume creation to fail.
- You cannot selecting RAID-5 for mirroring.
- Selecting "stripe across enclosures" is not recommended because then you need four enclosures, instead of two.
- Logging can slow performance.

Viewing the available disk storage

Before creating disk groups and volumes you may want to view available disk storage.

To view the available disk storage

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.

   On Windows 2012 operating systems, use the **Apps** menu.

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 disk group

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.

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

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

To create a dynamic disk group

1. Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, 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 as follows:
   - Enter the disk group name (for example, **DG1**).
   - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
   - 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.
   - Click **Next**.

   Note: Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).
7. Click **Next** to accept the confirmation screen with the selected disks.
8. Click **Finish** to create the dynamic disk group.

Adding disks to campus cluster sites

For campus cluster storage, Veritas recommends using Storage Foundation (SFW) site-aware allocation. To enable site-aware allocation, you assign a site name to

disks after they are added to a disk group. In the VEA assigning a site name is referred to as adding disks to a site.

For example, Disk1 and Disk2 are physically located on Site A and Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to site_a and add Disk3 and Disk4 to site_b.

To add disks to a site

1. From the VEA console, right-click a disk that needs to be added to a site and select **Add Disk to Site**.

   Disks must be part of a dynamic disk group in order to add them to a site.

2. In the Add Disk to a Site screen, choose one of the following:

   - Choose **Select a new site** and specify a new site name.
     The site name can include any alphanumeric value and valid characters like the period (.), dash (-), and underscore (_). It cannot exceed 31 characters. Site names are case insensitive; all names are converted to lowercase.
   - Choose **Available Sites** and select a site from the list.

3. From the **Available Disks** column, select the disk or disks to add to the specified site.

4. Click **OK**.

Creating volumes for campus clusters

This section will guide you through the process of creating a volume on a dynamic disk group for a campus cluster.

For creating volumes for other types of clusters:

- See "Creating dynamic volumes" on page 475.

Use the following procedure to create dynamic volumes for a campus cluster.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

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

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 as follows:

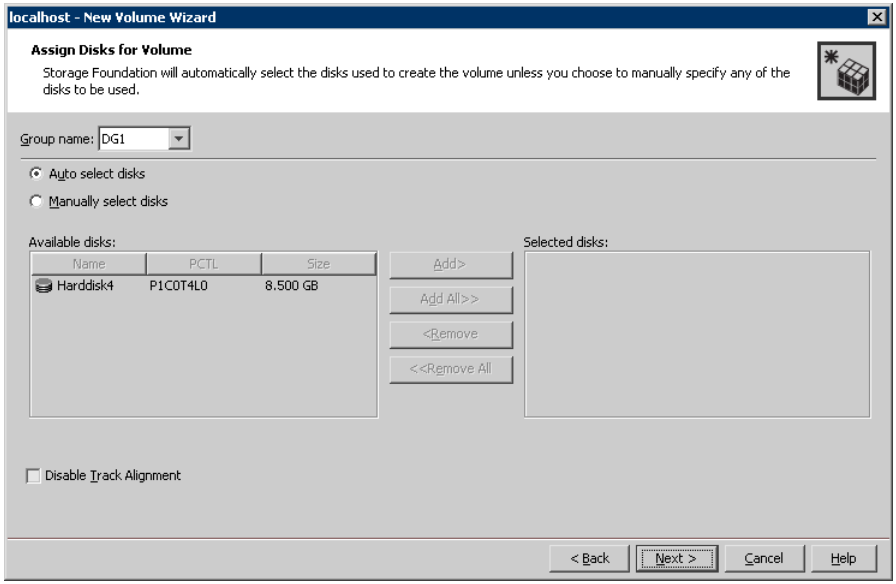

Group name	Make sure the appropriate disk group is selected.
Site preference	Select the **Site Separated** option.
Select site from	Select the campus cluster sites. Press **CTRL** to select multiple sites. **Note:** If no sites are listed, the disks have not yet been added to a site.

Auto select disks	Automatic disk selection 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 picks two disks (one from each array) with the most space.
Manually select disks	If you manually select disks, use the **Add** and **Remove** buttons to move the appropriate disks to the **Selected disks** list.
Disable Track Alignment	You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling track alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

Click **Next**.

7 Specify the volume attributes as follows:

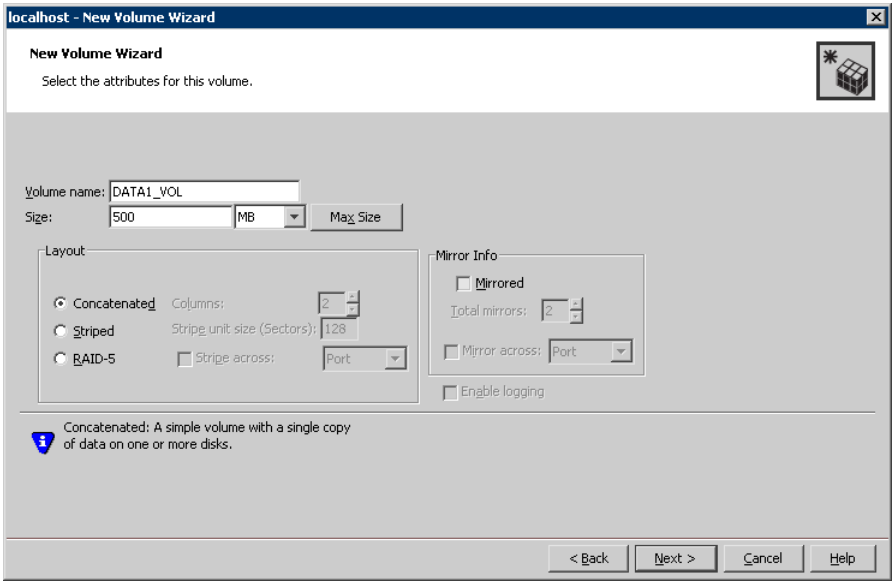

Volume name	Specify a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
Size	Specify a size for the volume. If you click **Max Size**, the **Size** box shows the maximum possible volume size for that layout in the dynamic disk group.
Layout	Ensure that the **Mirrored** checkbox is selected. Select either the **Concatenated** or **Striped** layout type. 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.
Mirror Info	Click **Mirror across** and select **Enclosures** from the drop-down list. When creating a site separated volume, as required for campus clusters, the number of mirrors must correspond to the number of sites. If needed, you can add more mirrors after creating the volume.
Enable logging	Verify that this option is not selected.

Click **Next**.

8 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.

Click **Next**.

9 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**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes as needed.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

Implementing a dynamic quorum resource

One of the key advantages of using SFW with Microsoft clustering is that you can create a mirrored quorum resource that adds fault tolerance to the quorum.

For information about tasks for creating a mirrored quorum resource, refer to the following topics:

- See "Creating a dynamic cluster disk group and a mirrored volume for the quorum resource" on page 515.

- See "Adding a VMDg resource for the quorum" on page 515.
- See "Changing the quorum resource to a dynamic mirrored quorum resource" on page 516.

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. Veritas 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 504.

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 a VMDg resource for the quorum

You 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

1. If Failover Cluster Management is already open, then proceed to the next step.

   To launch Failover Cluster Management, select it from **Start > Administrative Tools > Failover Cluster Management**.

   On Windows 2012 operating systems, use the **Start** screen.
2. Verify that the cluster is online on the same node where you created the disk group.

3. In the left pane of Failover Cluster Management, right-click **Services and Applications** and select **More Actions > Create Empty Service or Application**.
4. Right-click the new group and rename it, for example **QUORUM**.
5. Right-click **QUORUM** and select **Add a resource > More resources > Add Volume Manager Disk Group**.
6. Right-click **New Volume Manager Disk Group** in the center pane and click **Properties**.
7. In the General tab of the Properties dialog box, type a name for the resource in the **Resource Name** field, for example, **QUORUM_DG_RES**.
8. On the Properties tab, in the **Disk Group Name** field, type the name of the disk group that you previously created for the quorum, and click **OK** to close the dialog box.
9. Right-click the Quorum disk group resource (for example, QUORUM_DG_RES) in the left pane and select **Bring this resource online**.

   The specified disk group resource, QUORUM_DG_RES resource, is created under the Quorum group (for example, QUORUM).

Changing the quorum resource to a dynamic mirrored quorum resource

Use the following procedure to configure the cluster quorum settings and change the quorum resource to a dynamic mirrored quorum resource.

To change the quorum to a dynamic mirrored quorum resource

1. In Failover Cluster Management, right-click the cluster node in the configuration tree, and select **More Actions > Configure Cluster Quorum Settings**.
2. The Configure Cluster Quorum Wizard opens. Review the screen and click **Next**.
3. Select either the **Node and Disk Majority** or No **Majority: Disk Only** radio button, and click **Next**.
4. Select the storage resource that you want to assign as the disk witness for the quorum and click **Next**.

   This is the Volume Manager Disk Group resource that you previously created for the quorum disk group, for example, **QUORUM_DG_RES**.
5. Review the information in the Confirmation screen and click **Next**.
6. Click **Finish** to close the wizard.

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 519.

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.

   On Windows 2012 operating systems, use the **Start** screen.

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**.

   On Windows 2012 operating systems, use the **Start** screen.

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 18-4 Application group dependencies

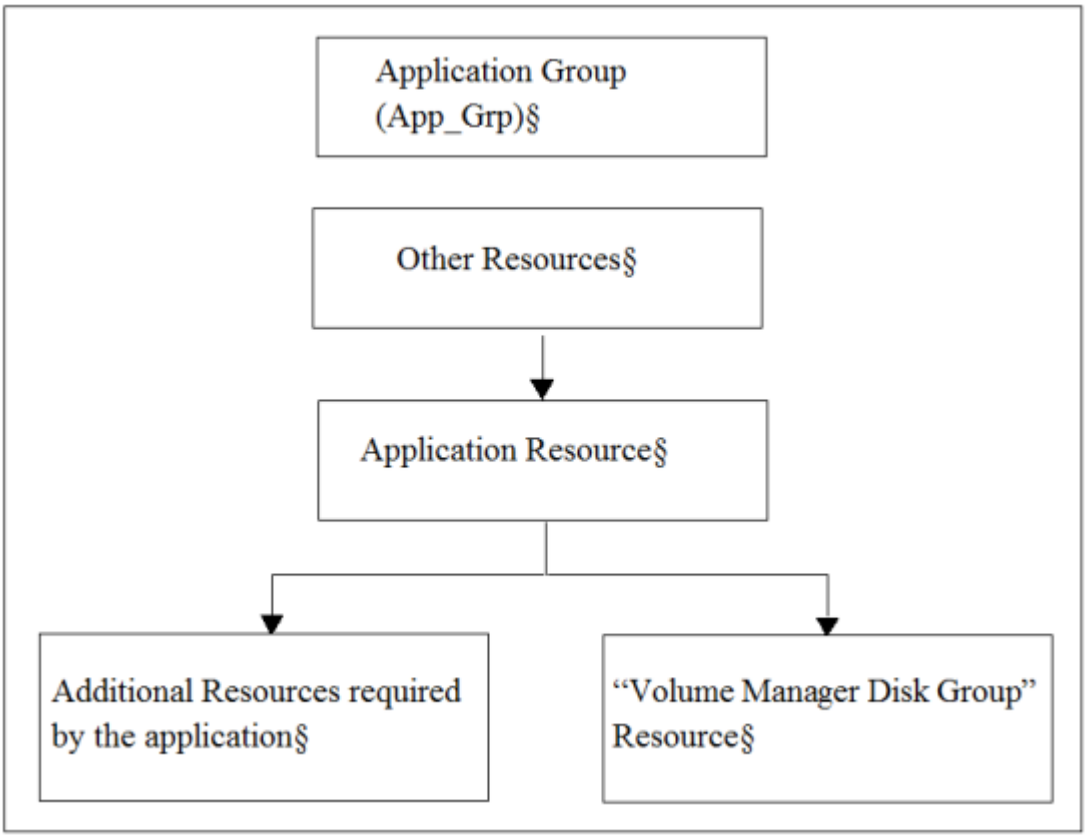

- Testing: After the application group is set up, test it by moving the cluster resources to another node and then move them back.

Verifying the cluster configuration

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, as follows:

- 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.

Caution: Do not simulate failover in a production environment.

To move online groups

1. Open the Failover Cluster Management tool by clicking **Start > Administrative Tools > Failover Cluster Management**.

   On Windows 2012 operating systems, use the **Start** screen.

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 *nameOfNode***.

   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 > Administrative Tools > Failover Cluster Management** from any node in the cluster.

   On Windows 2012 operating systems, use the **Start** screen.

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 the first step, select the resource group, and use **Move this service or application to another node > Move to node *nameOfNode*** to move the resource group.

Chapter 19

Deploying SFW and VVR with Microsoft failover clustering

This chapter includes the following topics:

- Tasks for deploying InfoScale Storage and Volume Replicator with Microsoft failover clustering
- Part 1: Setting up the cluster on the primary site
- Part 2: Setting up the cluster on the secondary site
- Part 3: Adding the Volume Replicator components for replication
- Part 4: Maintaining normal operations and recovery procedures

Tasks for deploying InfoScale Storage and Volume Replicator with Microsoft failover clustering

You can set up a disaster recovery (DR) solution using SFW with a Microsoft failover cluster and Volume Replicator on Windows Server systems.

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-Volume Replicator disaster recovery solution has four main parts:

- See "Part 1: Setting up the cluster on the primary site" on page 526.
- See "Part 2: Setting up the cluster on the secondary site" on page 535.

- See "Part 3: Adding the Volume Replicator components for replication" on page 537.
- See "Part 4: Maintaining normal operations and recovery procedures" on page 555.

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 "Tasks for deploying InfoScale Storage with Microsoft failover clustering" on page 458.

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 Volume Replicator option is selected during the InfoScale Storage installation and to configure the 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, Volume Replicator 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 19-1 Tasks for deploying InfoScale Storage with Microsoft failover clustering and Volume Replicator

Objective	Tasks
See "Part 1: Setting up the cluster on the primary site" on page 526.	
See "Reviewing the prerequisites and the configuration" on page 526.	■ Verify hardware and software prerequisites. ■ Review configuration requirements.
See "Installing and configuring the hardware" on page 529.	Set up and configure the hardware according to the manufacturers' instructions.
See "Installing Windows and configuring network settings" on page 529.	■ Install the operating system on both nodes. ■ Make necessary networking settings on both nodes.
See "Establishing a Microsoft failover cluster" on page 529.	Refer to Microsoft documentation for instructions on establishing a Microsoft failover cluster.

Table 19-1 Tasks for deploying InfoScale Storage with Microsoft failover clustering and Volume Replicator *(continued)*

Objective	Tasks
See "Installing InfoScale Storage (primary site)" on page 531.	See *Veritas InfoScale Installation and Upgrade Guide*.
See "Setting up security for Volume Replicator" on page 531.	Complete the steps to configure VxSAS.
See "Creating SFW disk groups and volumes" on page 534.	■ 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.
See "Creating a group for the application in the failover cluster" on page 478.	■ 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.
See "Installing the application on cluster nodes" on page 479.	■ 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.
See "Completing the setup of the application group in the failover cluster" on page 480.	■ 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 19-1 Tasks for deploying InfoScale Storage with Microsoft failover clustering and Volume Replicator *(continued)*

Objective	Tasks
See "Implementing a dynamic quorum resource" on page 481.	■ 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.
See "Verifying the cluster configuration" on page 520.	■ 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.
See "Completing the primary site configuration" on page 535.	Complete the primary site configuration.
See "Part 2: Setting up the cluster on the secondary site" on page 535.	
See "Repeating cluster configuration steps for the secondary site" on page 536.	■ Install and configure hardware ■ Install Windows and configure network settings ■ Establish the Microsoft failover cluster ■ Install SFW with the Volume Replicator option and Microsoft Cluster option ■ Install Volume Replicator Security Services (VxSAS) ■ Create SFW disk groups and volumes ■ Set up a group for the application in Failover Cluster Management ■ Install the application on cluster nodes ■ Complete the setup of the application group in Failover Cluster Management ■ Change the quorum resource to the dynamic quorum resource ■ Perform a final testing of the cluster
See "Part 3: Adding the Volume Replicator components for replication" on page 537.	

Table 19-1 Tasks for deploying InfoScale Storage with Microsoft failover clustering and Volume Replicator *(continued)*

Objective	Tasks
See "Configuring the Replicator Log volumes for Volume Replicator" on page 538.	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.
See "Creating the Replicated Data Sets with the wizard" on page 539.	Create Replicated Data Sets with Volume Replicator's Replicated Data Set wizard and start replication for the primary and secondary sites.
See "Creating resources for Volume Replicator" on page 552.	In Failover Cluster Management, create the network name and IP resource to be used for Volume Replicator replication.
See "Creating an RVG resource and setting the dependencies" on page 552.	■ 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.
See "Part 4: Maintaining normal operations and recovery procedures" on page 555.	
See "Normal operations: Monitoring the status of the replication" on page 556.	■ Monitor replication. ■ Perform planned migration.
See "Disaster recovery procedures" on page 556.	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.

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, refer to the illustration in the topic:

See “About a disaster recovery solution” on page 316.

Figure 19-1 DR configuration primary site

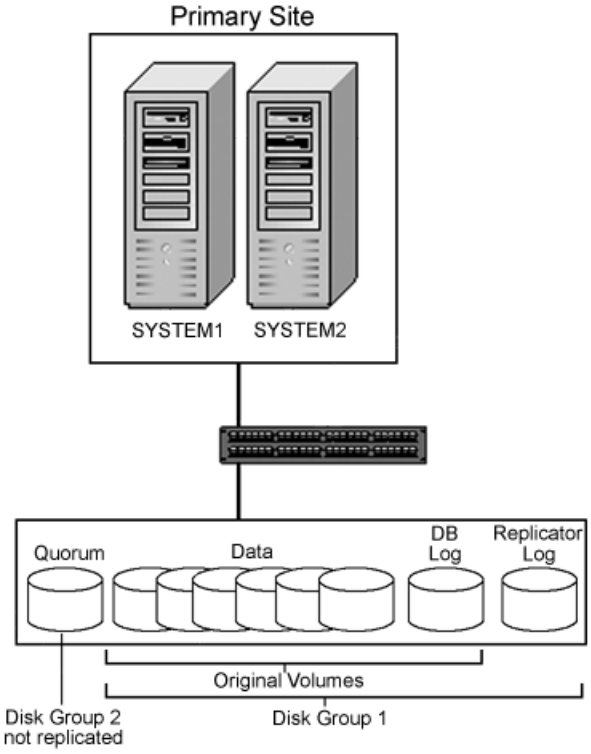

This configuration does not include DMP. For information about DMP and clustering:

See "Overview of configuration tasks for adding DMP DSMs" on page 178.

The following are some other key points about the configuration:

- A Microsoft failover cluster must be running before you install InfoScale Storage. Installing InfoScale Storage requires a reboot, but a reboot on the active cluster node causes it to fail over. Thus, Veritas recommends that you use a "rolling install" procedure to install InfoScale Storage 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, Veritas 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 InfoScale Storage 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 a Microsoft failover cluster

Before installing InfoScale Storage, you must first verify that Microsoft failover clustering is enabled (on a new Windows Server installation), 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**).

   On Windows 2012 operating systems, launch **Administrative Tools** from the **Start** screen.
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. Veritas 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 On 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 InfoScale Storage (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 Volume Replicator option from the product installer Options screen.

See "Tasks for installing InfoScale Foundation or InfoScale Storage for Microsoft failover clustering" on page 502.

Setting up security for Volume Replicator

If you use Volume Replicator for replication, you must configure the Veritas Volume Replicator Security Service (VxSAS) on all the cluster nodes.

In a Replicated Data Cluster environment, you must configure the service on all the nodes in the primary zone as well as the secondary zone.

For details on this required service, see the *Volume Replicator Administrator's Guide*.

After you install InfoScale Storage or InfoScale Enterprise, launch the Veritas Volume Replicator Security Service Configuration Wizard. This wizard lets you complete the Volume Replicator security service configuration.

To do so, launch the wizard after you install InfoScale Enterprise on both the primary and secondary nodes. Then, when you run the wizard, you can specify the primary and secondary sites in one step.

Prerequisites for configuring VxSAS

- The wizard requires you to be logged on with administrative privileges.
- The account that 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.
- The systems on which you want to configure VxSAS must be accessible from the local system.

To configure VxSAS

1. Launch the Veritas Volume replicator Security Service Configuration Wizard from **Start > All Programs > Veritas > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

   Optionally, run `vxsascfg.exe` from the command prompt of the required machine.

2. Read the information provided on the Welcome page and click **Next**.

3. Complete the Account Information panel as follows:

Account name (domain\account)	Enter the administrative account name.
Password	Specify a password

   If you have already configured VxSAS for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring VxSAS on the other hosts.

   Click **Next**.

4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 lets you specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

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 VxSAS.

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.

When configuring VxSAS in a firewall setup, the VxSAS wizard may not be able to configure the machines that are across the firewall, although the Host Selection dialog may list these nodes. In this case, configure VxSAS locally on the machines that are across the firewall.

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 Volume Replicator 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. Veritas recommends using three disks for the mirrored quorum for additional redundancy.

Figure 19-2 Microsoft clustered database with disks for data, logs, and the quorum resource

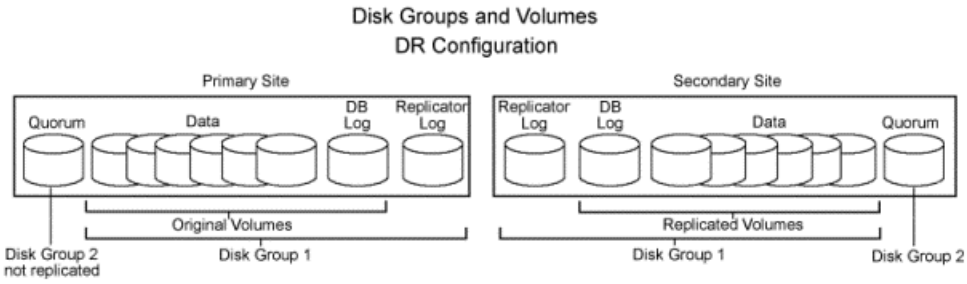

Do not use the following types of volumes for the data and Replicator Log volumes; Volume Replicator does not support these types of volumes:

- Storage Foundation (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, refer to the following topic:

See “Creating SFW disk groups and volumes” on page 472.

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; see the following topics in that chapter to complete configuring the primary site:

- See "Creating a group for the application in the failover cluster" on page 478.
- See "Installing the application on cluster nodes" on page 479.
- See "Completing the setup of the application group in the failover cluster" on page 480.
- See "Implementing a dynamic quorum resource" on page 481.
- See "Verifying the cluster configuration" on page 520.

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 Veritas recommends that the sizes be the same. Install the application on the secondary cluster nodes the same as on the primary cluster.

Figure 19-3 DR configuration secondary site

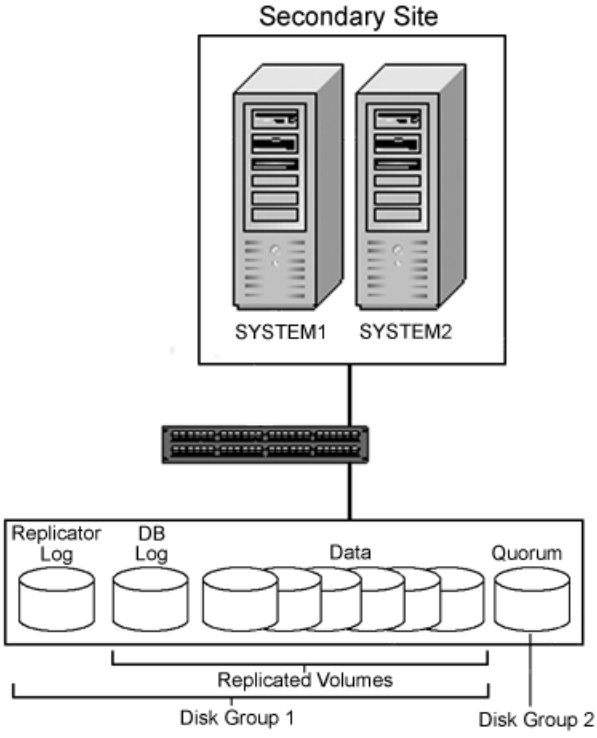

Repeating cluster configuration steps for the secondary site

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 InfoScale Storage and Volume Replicator with Microsoft failover clustering" on page 522.

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 Volume Replicator components to the configuration.

See "Part 3: Adding the Volume Replicator components for replication" on page 537.

Part 3: Adding the Volume Replicator components for replication

This section provides information on configuring the Volume Replicator components for replication. Topics include:

- See "Volume Replicator components overview" on page 537.
- See "Configuring the Replicator Log volumes for Volume Replicator" on page 538.
- See "Creating the Replicated Data Sets with the wizard" on page 539.
- See "Creating resources for Volume Replicator" on page 552.
- See "Creating an RVG resource and setting the dependencies" on page 552.

Volume Replicator components overview

You configure the following Volume Replicator components:

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, 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. 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 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. Veritas recommends having Replicator Log volumes of the same size at the primary site and the secondary site.

Configuring the Replicator Log volumes for Volume Replicator

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, Veritas 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 Volume Replicator

1. Click **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop to open the VEA console on the active node of the primary site.

   On Windows 2012 operating systems, use the **Apps** menu.

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 *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.

Creating the Replicated Data Sets with the wizard

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.

Configuring Volume Replicator 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.

If you are using IPv6, make sure that you have set the appropriate IP preference for Volume Replicator before configuring Volume Replicator replication. The default setting is IPv4. When you specify host names while configuring replication, Volume

Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP protocol Volume Replicator uses to resolve the host names. Use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to set the IP preference.

- Verify that the data volumes are not of the following types, as Volume Replicator does not support these types of volumes:
  - Storage Foundation (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 disk group is imported and the volumes are mounted in the primary and secondary zone.
- Verify that you have set the appropriate IP preference, whether Volume Replicator should use IPv4 or IPv6 addresses. The default setting is IPv4. When you specify host names while configuring replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP protocol Volume Replicator uses to resolve the host names. Use Veritas Enterprise Administrator (VEA) (Control Panel > VVR Configuration > IP Settings tab) to set the IP preference.
- Configure the VxSAS service if you have not already done so.
  See "Setting up security for Volume Replicator" on page 531.

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.

   Start VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**.

   On Windows 2012 operating systems, from the **Apps** menu in the Start screen.

   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**.

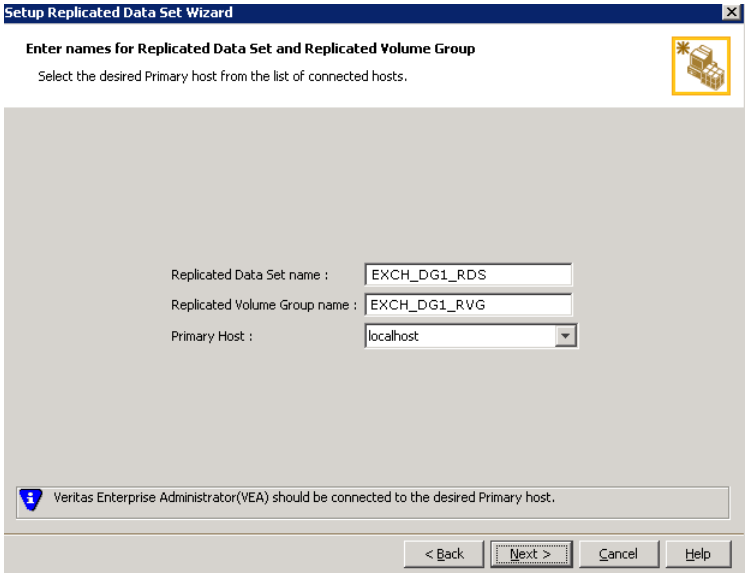

By default, the local host is selected as the Primary Host. To specify a different host name, make sure the required host is connected to the VEA console and select it in the Primary Host list.

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**.

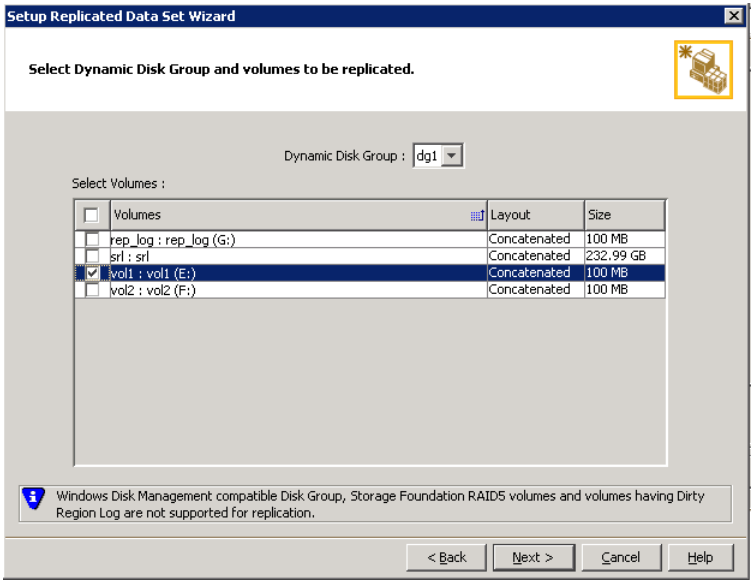

To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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:

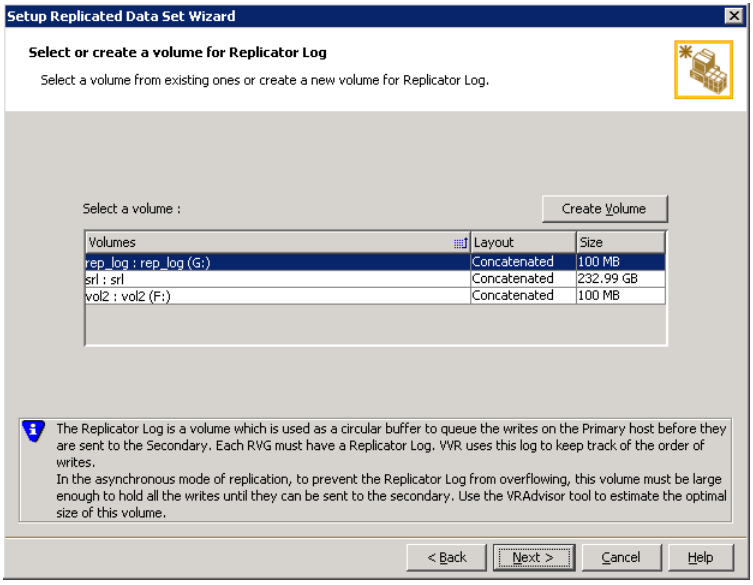

To select an existing volume

- 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 appears:

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 *Storage Foundation Administrator's Guide*.
- Choose the **Select disks automatically** option if you want Volume Replicator 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, Volume Replicator 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 Volume Replicator displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

- The same or larger amount of space as that on the Primary
- Enough space to create volumes with the same layout as on the Primary Otherwise, the RDS setup wizard enables you to create the required volumes manually.
- Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
- Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.

11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

- If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

  Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log.

  When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.

12 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

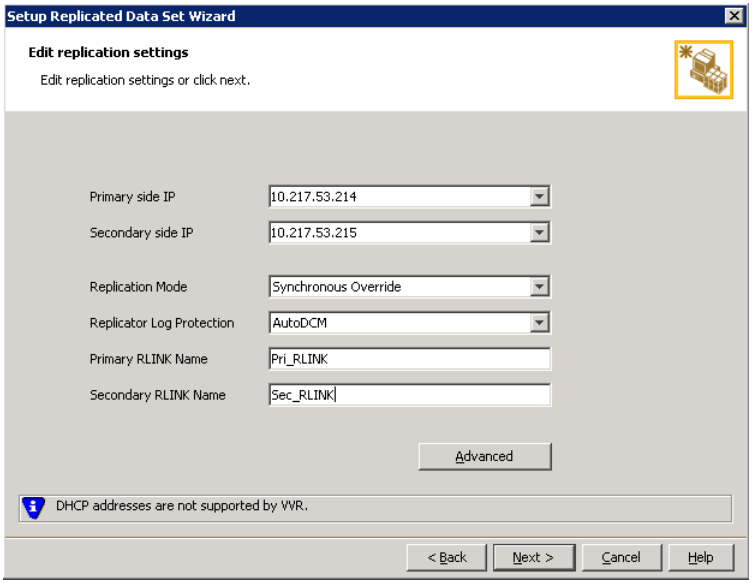

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side	IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.
Secondary side IP	Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode

Select the required mode of replication:

- **Synchronous Override** (default) 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 Volume Replicator updates in the Replicator Log. From there, Volume Replicator 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	■ **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 Volume Replicator 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 Volume Replicator assigns a default name.

- If you want to specify advanced replication settings, click **Advanced**. Edit the replication settings for a secondary host as needed.

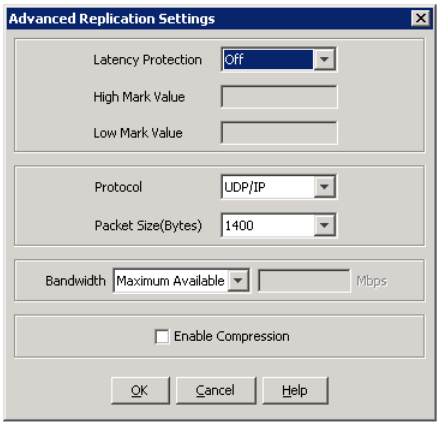

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

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, Volume Replicator 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.
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, Volume Replicator 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**.

13 On the Start Replication page, choose the appropriate option as follows:

- To add the Secondary and start replication immediately, select **Start Replication** with one of the following options:

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, Volume Replicator 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 and ReFS 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 *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.

14 Review the information.

Click **Back** to change any information you had specified.

Otherwise, 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 this procedure for each additional disk group. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name.

Creating resources for Volume Replicator

Create the resources for Volume Replicator 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 Volume Replicator replication.

A separate valid IP address is necessary for Volume Replicator replication, because on the secondary cluster before a disaster, the application IP must be offline whereas the Volume Replicator 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 Volume Replicator 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:

- See "Creating an RVG resource" on page 553.
- See "Setting the application resource dependency on the RVG resource" on page 553.

Creating an 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**.

   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 **Click here to add a dependency** box.
   - From the **Resource** drop-down list, select the network name you created for the RVG. Click **Insert**.
   - Click the **Click here to add a dependency**box.
   - 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**.
   - Click **OK** to close the Properties dialog box.
4. 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 19-4 Dependencies of Volume Replicator-related resources

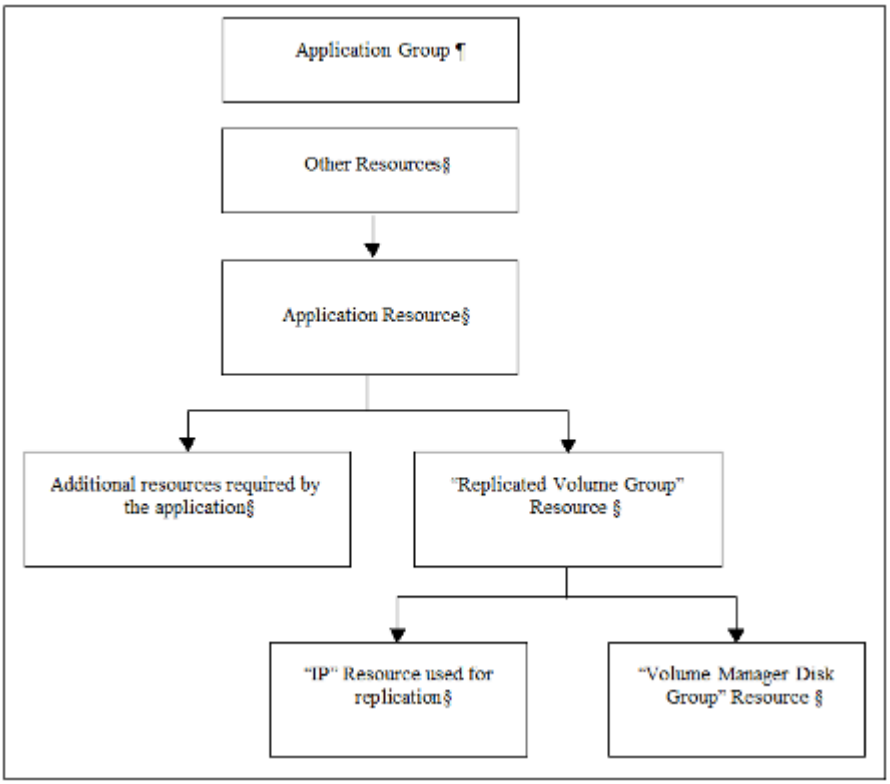

The chart shows only the Volume Replicator-related resources. Normally, there would be other resources involved in any clustered application. The main point of the chart is to show that the RVG resource is now dependent on the Volume Manager Disk Group resource and the Volume Replicator virtual IP resource. 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 *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 the following procedure.

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.

   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.

3. Make sure that the volumes have retained the same drive letters as they had before the disaster.

4. Bring the application group online on the original Primary.

Section 8

Server Consolidation

- Chapter 20. Server consolidation overview
- Chapter 21. Server consolidation configurations

Chapter 20

Server consolidation overview

This chapter includes the following topics:

- Server consolidation definition
- Need for implementing server consolidation
- Advantages of using SFW with server consolidation
- Overview of the server consolidation process

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 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 supports clustering with MSCS and Storage Foundation and High Availability Solutions includes 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.

Chapter 21

Server consolidation configurations

This chapter includes the following topics:

- Typical server consolidation configuration
- Server consolidation configuration 1—many to one
- Server consolidation configuration 2—many to two: Adding clustering and DMP
- SFW features that support server consolidation

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 21-1 General server consolidation configuration

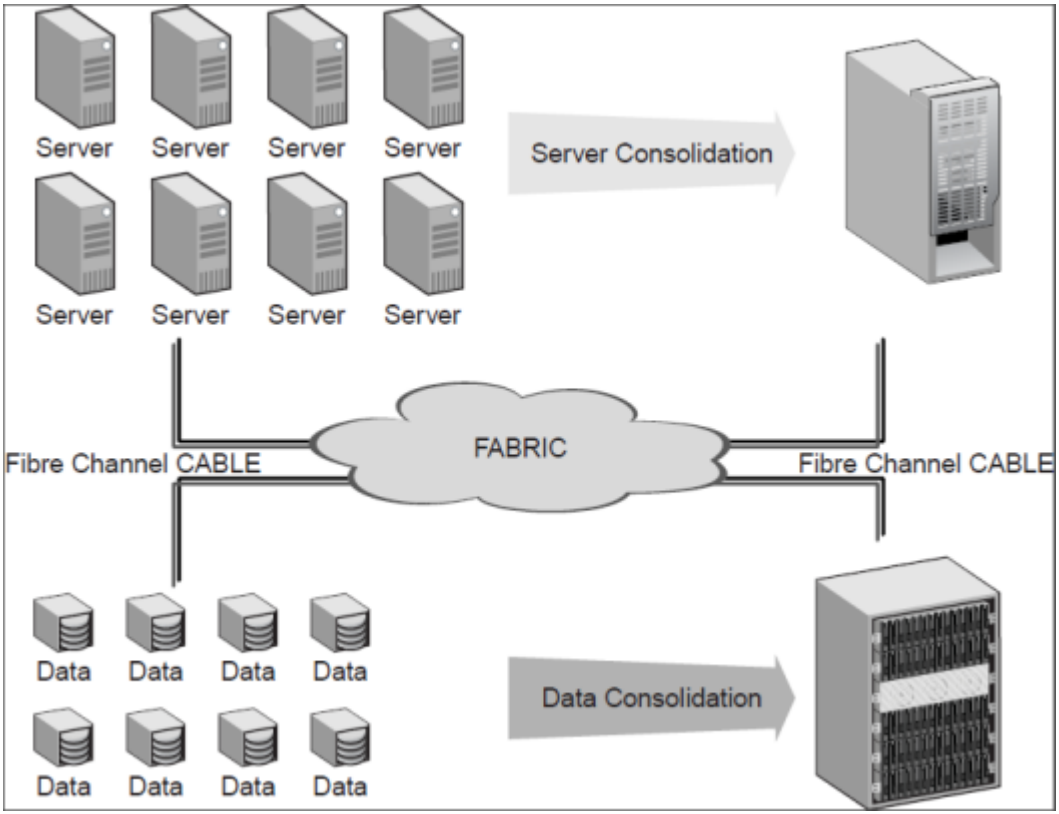

Proof of concept

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 21-2 Proof of concept: Consolidating four small servers to one large server

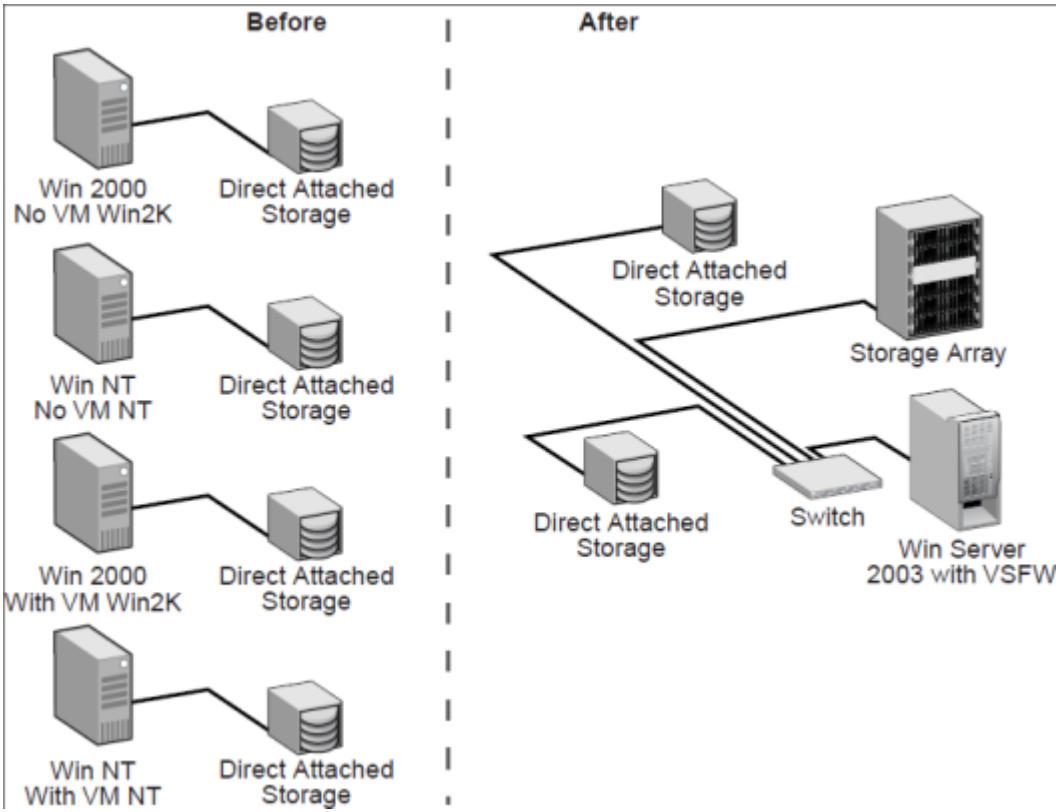

About this configuration

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 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 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 21-1 Tasks for server consolidation for many to one configuration

Objectives	Tasks
See "Preparing to consolidate" on page 569.	■ 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 the InfoScale product. 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.
See "Migrating the data to the large server" on page 570.	■ 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.
See "Adding the storage array" on page 571.	■ 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.
See "Completing the consolidation process" on page 572.	■ Migrate the data from the remaining servers.

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.

   Warning: 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 operating system and InfoScale product 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 (no VM or InfoScale product 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 2012 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 2012 features in SFW. Use the Upgrade Dynamic Disk Group Version command.

7 Test the data on the Windows Server 2012 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 2012 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 Microsoft clustering and DMP.

Figure 21-3 Adding fault tolerance with Microsoft clustering and DMP — requires two servers

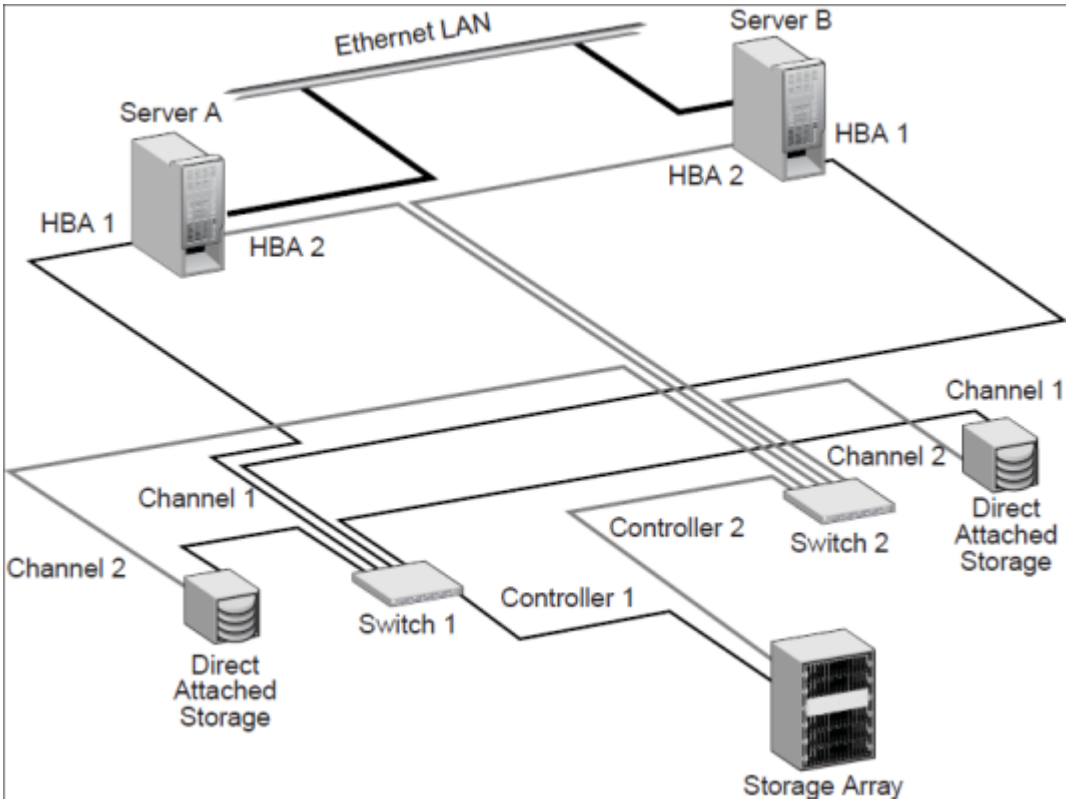

About this configuration

This configuration is an upgrade to Server Consolidation Configuration 1, to add Microsoft clustering 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 21-2 Tasks for server consolidation adding Microsoft clustering and DMP

Objectives	Tasks
See "Adding the new hardware" on page 575.	■ 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.
See "Establishing the Microsoft failover cluster" on page 576.	■ Refer to Microsoft instructions for establishing the cluster under Microsoft clustering.
See "Adding SFW support to the cluster" on page 576.	■ With Server B as the active cluster node, use **Add or Remove Programs** to add DMP and the Microsoft clustering option to the first server. ■ With Server A as the active node, install InfoScale Storage with the DMP and Microsoft clustering options to Server B. ■ Change the existing disk groups to cluster disk groups. ■ Prepare a disk group for the dynamic mirrored quorum.
See "Setting up Microsoft failover cluster groups for the applications" on page 577.	■ If you have applications on the server that you want to cluster, create Microsoft failover cluster groups for them.
See "Installing applications on the second computer" on page 578.	■ Install the applications' program files on the local drive of Server B.
See "Completing the setup of the application group in the Microsoft cluster" on page 578.	■ Complete the cluster application group by adding resources and setting dependencies.
See "Changing the quorum resource to the dynamic quorum resource" on page 578.	■ 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.

Table 21-2 Tasks for server consolidation adding Microsoft clustering and DMP *(continued)*

Objectives	Tasks
See "Verifying the cluster configuration" on page 578.	■ Test the cluster by moving the cluster resources to the other node.
See "Enabling DMP" on page 578.	■ 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.

Warning: 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.

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.

   Warning: 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 Microsoft failover cluster

Complete the steps necessary to install a cluster on Server A and Server B, using Microsoft clustering. 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.
- Install the Windows Server 2012 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 Microsoft clustering 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 InfoScale Storage with the DSMs and Microsoft clustering 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 Microsoft failover cluster groups for the applications

If you have applications on the server that you want to cluster, you need to set up a 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 Microsoft failover cluster groups:

See "Creating a group for the application in the failover cluster" on page 478.

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 Microsoft clustering environment:

See "Installing the application on the cluster nodes" on page 518.

Completing the setup of the application group in the Microsoft cluster

Once the application is installed, complete the configuration of the application group in the Microsoft cluster. For details:

See "Completing the setup of the application group in the cluster" on page 519.

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 514.

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 520.

Enabling DMP

These steps assume that InfoScale Storage with the DMP DSMs has been installed.

See "Adding SFW support to the cluster" on page 576.

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:
   - 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.

- Uncheck the **Exclude** checkbox.

2 Using appropriate cables, connect the second path on Server A to Switch 2.

- Connect the path through Server A, HBA 2, Channel 2 of the direct-attached storage, and Controller 2 of the large storage array.
- 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 DSMs is installed correctly.

4 Complete step 1 to step 3 on Server B.

Microsoft clustering and DMP are now set up, and the upgraded configuration steps are complete.

SFW features that support server consolidation

With consolidated servers, Storage Foundation has multiple features that assure fault tolerance and improve storage utilization.

See "Advantages of using SFW with server consolidation" on page 562.

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:

See "Automatic volume growth" on page 579.

See "Features that support storage in a SAN" on page 580.

See "Performance monitoring" on page 580.

Automatic volume growth

Storage Foundation 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 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 Storage Foundation Administrator's Guide.