

Veritas Storage Foundation™ and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL

Windows Server 2003,
Windows Server 2008

5.1



SFW HA and Disaster Recovery Solutions Guide for Microsoft SQL

Copyright © 2008 Symantec Corporation. All rights reserved.

Storage Foundation 5.1 for Windows HA

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

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

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

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Windows is a registered trademark of Microsoft Corporation.

Licensing and registration

Storage Foundation for Windows and Storage Foundation HA for Windows are licensed products. See the *Veritas Storage Foundation and High Availability Solutions for Windows, Installation and Upgrade Guide* for license installation instructions.

Technical support

For technical assistance, visit

<http://www.symantec.com/business/support/index.jsp> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Section 1 Introduction

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

About the solutions guides	26
About high availability	26
About campus clusters	27
About replicated data clusters	27
About disaster recovery	27
How this guide is organized	28

Chapter 2 Using the Solutions Configuration Center

About the Solutions Configuration Center	29
Starting the Configuration Center	30
Available options from the Configuration Center	30
About running the Configuration Center wizards	37
Following the workflow in the Configuration Center	38
Solutions wizard logs	41

Section 2 High Availability

Chapter 3 High availability for SQL: Overview

What is high availability?	45
Why implement a high availability solution?	46
How the agent makes SQL Server highly available	46

Chapter 4 Deploying SFW HA for high availability: New SQL Server 2000 installation

Tasks for a new HA installation of SQL Server 2000	48
Tasks for an Active-Passive configuration	49
Tasks for an Active-Active configuration	51
Reviewing the requirements	53
Disk space requirements	53

Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	53
Reviewing the configuration	57
Active-Passive configuration	58
Active-Active configuration	60
Disaster recovery configuration	62
Reviewing considerations for Active-Active configurations	65
Key information for Active-Active configurations	66
Following the workflow in the Solutions Configuration Center	67
Configuring the storage hardware and network	68
Installing Veritas Storage Foundation HA for Windows	71
Setting Windows driver signing options	71
Installing Storage Foundation HA for Windows	73
Configuring cluster disk groups and volumes for SQL Server 2000	77
About cluster disk groups and volumes	77
Prerequisites for configuring cluster disk groups and volumes	78
Sample disk group and volume configuration	78
Considerations for disk groups and volumes for multiple instances	79
Considerations for volumes for a disaster recovery configuration	80
Creating a cluster disk group	81
Creating volumes	82
Configuring the cluster	86
Configuring Web console	98
Configuring notification	99
About installing multiple instances	102
Installing and configuring SQL Server 2000 on the first node	103
Installing Microsoft SQL Server 2000	103
Setting SQL Server 2000 services to manual start	106
Preparing to install SQL Server 2000 on the second node	106
Stopping the SQL Server 2000 service	107
Deporting the cluster disk group	107
Importing the cluster disk group	108
Adding drive letters to mount the volumes	108
Renaming shared SQL Server 2000 files	110
Installing and configuring SQL Server 2000 on the second node	110
Installing SQL Server 2000 on the second node	110
Removing shared SQL Server 2000 files	113
Setting the internal name of the clustered instance	114
Creating a SQL Server user-defined database	116
Completing configuration steps in SQL Server	117

Moving the tempdb database if using VVR for disaster recovery	117
Assigning ports for multiple SQL Server instances	118
Configuring the VCS SQL Server 2000 service group	119
Service group requirements for active-active configurations	119
Prerequisites for configuring the service group	120
Creating the SQL Server 2000 service group	121
Verifying the SQL Server 2000 cluster configuration	125
Determining additional steps needed	126
Configuring the Cluster Management Console connection	127
Prerequisites for installing the cluster connector	127
Installing the cluster connector on Windows clusters	128
Configuring the cluster connector	129
Modifying the SQL 2000 service group to add VMDg and MountV resources	131

Chapter 5

Deploying SFW HA for high availability: Standalone SQL Servers

Tasks for converting a standalone SQL 2000 Server for high availability	134
Reviewing the requirements	137
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	137
Reviewing the configuration	141
Sample configuration	142
Configuring the storage hardware and network	143
Preparing the standalone SQL Server	146
Installing Veritas Storage Foundation HA for Windows	147
Setting Windows driver signing options	147
Installing Storage Foundation HA for Windows	149
Configuring cluster disk groups and volumes for SQL Server 2000	153
About cluster disk groups and volumes	154
Prerequisites for configuring cluster disk groups and volumes	154
Sample disk group and volume configuration	155
Considerations for converting existing shared storage to cluster disk groups and volumes	156
Considerations for volumes for a VVR disaster recovery configuration	156
Creating a cluster disk group	157
Creating volumes	159

Configuring the cluster	163
Configuring Web console	175
Configuring notification	176
Verifying that SQL Server 2000 databases and logs are moved to shared storage	179
Preparing to install SQL Server 2000 on additional nodes	179
Deporting the cluster disk group	180
Importing the cluster disk group	180
Adding drive letters to mount the volumes	181
Renaming shared SQL Server 2000 files	182
Installing and configuring SQL Server 2000 on additional nodes	182
Installing SQL Server 2000 on additional nodes	182
Setting the internal name of the clustered instance	186
Completing configuration steps in SQL Server	188
Moving the tempdb database if using VVR for disaster recovery	188
Assigning the port for multiple instances	189
Configuring the VCS SQL Server 2000 service group	189
Prerequisites for configuring the service group	189
Creating the SQL Server 2000 service group	191
Verifying the SQL Server 2000 cluster configuration	195
Adding a SQL Server user-defined database	196
Additional instructions for disaster recovery	196

Chapter 6

Deploying SFW HA for high availability: New SQL Server 2005 installation

Tasks for a new HA installation of SQL Server 2005	198
Tasks for an Active-Passive configuration	199
Tasks for an Active-Active configuration	201
Reviewing the requirements	203
Disk space requirements	203
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	203
Reviewing the configuration	207
Active-Passive configuration	208
Active-Active configuration	210
Disaster recovery configuration	212
Reviewing considerations for Active-Active configurations	215
Key information for Active-Active configurations	216
Following the workflow in the Solutions Configuration Center	217
Configuring the storage hardware and network	218
Installing Veritas Storage Foundation HA for Windows	221

Setting Windows driver signing options	221
Installing Storage Foundation HA for Windows	223
Resetting the driver signing options	227
Configuring cluster disk groups and volumes for	
SQL Server 2005	228
About cluster disk groups and volumes	228
Prerequisites for configuring cluster disk groups and volumes	229
Sample disk group and volume configuration	229
Considerations for disk groups and volumes for multiple instances	230
Considerations for volumes for a disaster recovery configuration	230
Creating a cluster disk group	231
Creating volumes	233
Configuring the cluster	237
Configuring Web console	249
Configuring notification	250
About installing multiple instances	254
Installing and configuring SQL Server 2005 on the first node	254
Installing Microsoft SQL Server 2005 on the first node	255
Setting the startup mode of the SQL Server 2005 services	258
Preparing to install SQL Server 2005 on the second node	259
Stopping the SQL Server 2005 service	259
Deporting the cluster disk group	260
Importing the cluster disk group	260
Adding drive letters to mount the volumes	261
Renaming shared SQL Server 2005 files	262
Installing and configuring SQL Server 2005 on the second node	262
Installing Microsoft SQL Server on the second node	262
Removing shared SQL Server files	266
Setting the internal name of the clustered instance	266
Creating a SQL Server user-defined database	269
Completing configuration steps in SQL Server	270
Moving the tempdb database if using VVR for disaster recovery	270
Assigning ports for multiple SQL Server instances	271
Configuring the VCS SQL Server 2005 service group	272
Service group requirements for active-active configurations	272
Prerequisites for configuring the service group	273
Creating the SQL Server 2005 service group	274
Verifying the SQL Server 2005 cluster configuration	278

Determining additional steps needed	279
Configuring the Cluster Management Console connection	280
Prerequisites for installing the cluster connector	280
Installing the cluster connector on Windows clusters	281
Configuring the cluster connector	282
Modifying a SQL 2005 service group to add VMDg and MountV resources	284

Chapter 7

Deploying SFW HA for high availability: Standalone SQL 2005 Servers

Tasks for converting a standalone SQL 2005 Server for high availability	288
Reviewing the requirements	291
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	291
Reviewing the configuration	296
Sample configuration	297
Configuring the storage hardware and network	298
Preparing the standalone SQL Server	300
Backing up existing SQL data	300
Setting SQL Server services to manual start	300
Installing Veritas Storage Foundation for Windows	302
Setting Windows driver signing options	302
Installing Storage Foundation HA for Windows	304
Configuring cluster disk groups and volumes for SQL Server 2005	308
About cluster disk groups and volumes	308
Prerequisites for configuring cluster disk groups and volumes	309
Sample disk group and volume configuration	310
Considerations for converting existing shared storage to cluster disk groups and volumes	310
Considerations for volumes for a VVR disaster recovery configuration	311
Creating a cluster disk group	312
Creating volumes	314
Configuring the cluster	318
Configuring Web console	330
Configuring notification	331
Verifying that SQL Server 2005 databases and logs are moved to shared storage	334
Preparing to install SQL Server 2005 on the second node	334
Deploying the cluster disk group	335

	Importing the cluster disk group	335
	Adding drive letters to mount the volumes	336
	Renaming shared SQL Server 2005 files	337
	Installing and configuring SQL Server 2005 on the second node	337
	Installing SQL Server 2005 on the second node	337
	Removing shared SQL Server files	341
	Setting the internal name of the clustered instance	341
	Configuring the VCS SQL Server 2005 service group	344
	Prerequisites for configuring the service group	344
	Creating the SQL Server 2005 service group	345
	Verifying the SQL Server 2005 cluster configuration	349
	Adding a new SQL Server user-defined database	351
	Additional instructions for disaster recovery	351
Chapter 8	Configuring an MSDTC service group for high availability	
	Tasks for configuring MSDTC for high availability	353
	Reviewing the prerequisites	355
	Reviewing the configuration	355
	Configuring cluster disk groups and volumes	358
	Creating a cluster disk group	359
	Creating volumes	360
	Mounting volumes used by the MSDTC service group	364
	Creating an MSDTC service group	365
	Creating an MSDTC client	367
Section 3	Campus Cluster	
Chapter 9	Campus cluster for SQL Server: Overview	
	What is a campus cluster?	372
	Differences between campus clusters and local clusters	372
	Sample Campus Cluster configuration	373
	Why implement a campus cluster?	374
	Campus cluster failover using the ForceImport attribute	374
	Reinstating faulted hardware	376

Chapter 10	Deploying SFW HA for Campus Cluster: New SQL Server 2000 Installation	
	Reviewing the requirements	383
	Disk space requirements	383
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	383
	Reviewing the configuration	388
	Following the workflow in the Solutions Configuration Center	389
	Configuring the storage hardware and network	389
	Installing Veritas Storage Foundation HA for Windows	392
	Setting Windows driver signing options	392
	Installing Storage Foundation HA for Windows	394
	Resetting the driver signing options	398
	Configuring the cluster	398
	Configuring Web console	411
	Configuring notification	412
	Configuring cluster disk groups and volumes for SQL Server 2000	416
	About cluster disk groups and volumes	416
	Prerequisites for configuring cluster disk groups and volumes	416
	Sample disk group and volume configuration	417
	Configuring the disks and volumes	418
	Creating a dynamic (cluster) disk group	419
	Creating a volume	421
	Managing disk groups and volumes	425
	Importing a disk group and mounting a volume	425
	Unmounting a volume and deporting a disk group	426
	Installing and configuring SQL Server 2000 on the first node	427
	Installing Microsoft SQL Server 2000	427
	Setting SQL Server 2000 services to manual start	430
	Preparing to install SQL Server 2000 on the second node	430
	Stopping the SQL Server 2000 service	431
	Deporting the cluster disk group	431
	Importing the cluster disk group	432
	Adding drive letters to mount the volumes	432
	Renaming shared SQL Server 2000 files	434
	Installing and configuring SQL Server 2000 on the second node	434
	Installing SQL Server 2000 on the second node	434
	Removing shared SQL Server 2000 files	437
	Setting the internal name of the clustered instance	438
	Creating a SQL Server user-defined database	440

Completing configuration steps in SQL Server	441
Moving the tempdb database if using VVR for disaster recovery	441
Configuring the SQL Server 2000 service group for VCS	442
Prerequisites for configuring the SQL Server 2000 service group	442
Creating the SQL Server 2000 service group	443
Modifying the IP resource in the SQL Server 2000 service group	448
Verifying the campus cluster: Switching the service group	449
Setting the ForceImport attribute to 1 after a site failure	449
Modifying the SQL 2000 service group to add VMDg and MountV resources	450

Chapter 11 Deploying SFW HA for Campus Cluster: New SQL Server 2005 Installation

Reviewing the requirements	457
Disk space requirements	457
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	457
Reviewing the configuration	462
Following the workflow in the Solutions Configuration Center	463
Configuring the storage hardware and network	463
Installing Veritas Storage Foundation HA for Windows	466
Setting Windows driver signing options	466
Installing Storage Foundation HA for Windows	468
Resetting the driver signing options	472
Configuring the cluster	472
Configuring Web console	485
Configuring notification	486
Configuring cluster disk groups and volumes for SQL Server 2005	490
About cluster disk groups and volumes	490
Prerequisites for configuring cluster disk groups and volumes	490
Sample disk group and volume configuration	491
Configuring the disks and volumes	492
Creating a dynamic (cluster) disk group	493
Creating a volume	495
Managing disk groups and volumes	499
Importing a disk group and mounting a volume	499
Unmounting a volume and deporting a disk group	500

Installing and configuring SQL Server 2005 on the first node	501
Installing Microsoft SQL Server 2005	501
Setting the startup mode of the SQL Server 2005 services	504
Preparing to install SQL Server 2005 on the second node	505
Stopping the SQL Server 2005 service	505
Deporting the cluster disk group	505
Importing the cluster disk group	506
Adding drive letters to mount the volumes	506
Renaming shared SQL Server 2005 files	508
Installing and configuring SQL Server 2005 on the second node	508
Installing SQL Server 2005 on the second node	508
Removing shared SQL Server 2005 files	512
Setting the internal name of the clustered instance	512
Creating a SQL Server user-defined database	515
Completing configuration steps in SQL Server	516
Moving the tempdb database if using VVR for disaster recovery	516
Configuring the SQL Server 2005 service group for VCS	517
Prerequisites for configuring the SQL Server 2005 service group	517
Creating the SQL Server 2005 service group	518
Modifying the IP resource in the SQL Server 2005 service group	523
Verifying the campus cluster: Switching the service group	524
Setting the ForceImport attribute to 1 after a site failure	524
Modifying the SQL 2005 service group to add VMDg and MountV resources	525

Section 4 Replicated Data Clusters

Chapter 12 About Replicated Data Clusters

About Replicated Data Clusters	529
How VCS Replicated Data Clusters work	531
Setting up a Replicated Data Cluster configuration	532
Setting up replication	532
Configuring the service groups	533
Migrating the service group	535

Chapter 13	Configuring Replicated Data Clusters for SQL 2000	
	Tasks for configuring Replicated Data Clusters for SQL 2000	538
	Reviewing the prerequisites	541
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	541
	Reviewing the configuration	545
	Sample configuration	546
	Configuring the storage hardware and network	547
	Installing Veritas Storage Foundation HA for Windows	550
	Setting Windows driver signing options	550
	Installing Storage Foundation HA for Windows	552
	Configuring VxSAS	556
	Resetting the driver signing options	558
	Configuring the cluster	558
	Configuring Web console	571
	Configuring notification	572
	Configuring cluster disk groups and volumes	575
	Creating a cluster disk group	576
	Creating volumes	578
	Installing and configuring SQL Server 2000 on the first node	584
	Installing Microsoft SQL Server	584
	Setting SQL Server 2000 services to manual start	587
	Preparing to install SQL Server on the second node	587
	Stopping the SQL Server 2000 service	588
	Deporting the cluster disk group	588
	Importing the cluster disk group	589
	Adding drive letters to mount the volumes	589
	Renaming shared SQL Server 2000 files	591
	Installing SQL Server 2000 on the second node	591
	Installing SQL Server	591
	Removing shared SQL Server files	594
	Setting the internal name of the clustered instance	595
	Configuring the VCS SQL Server service group	597
	Creating the primary system zone	602
	Registering the virtual server in the SQL Server	603
	Creating a SQL Server user-defined database	605
	Creating new volumes	605
	Creating a new SQL Server database	605
	Adding VMDg and MountV resources	606
	Verifying the installation in the primary zone	607
	Creating a parallel environment in the secondary zone	608
	Adding the systems in the secondary zone to the cluster	609

Setting up the Replicated Data Sets (RDS)	617
Configuring a hybrid RVG service group for replication	628
Creating the RVG service group	628
Configuring the RVG service group for RDC replication	629
Configuring the RVG Primary resources	639
Configuring the primary system zone for the RVG	642
Setting a dependency between the service groups	643
Adding the nodes from the secondary zone to the RDC	644
Adding the nodes from the secondary zone to the RVG service group	644
Configuring secondary zone nodes in the RVG service group	647
Configuring the IP resources for failover	647
Adding the nodes from the secondary zone to the SQL Server service group	649
Configuring the zones in the SQL Server service group	650
Verifying the RDC configuration	651
Bringing the service group online	651
Switching online nodes	651
Additional instructions for GCO disaster recovery	652

Chapter 14 Configuring Replicated Data Clusters for SQL 2005

Tasks for configuring Replicated Data Clusters for SQL 2005	654
Reviewing the prerequisites	657
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	657
Reviewing the configuration	662
Sample configuration	662
Configuring the storage hardware and network	664
Installing Veritas Storage Foundation HA for Windows	666
Setting Windows driver signing options	666
Installing Storage Foundation HA for Windows	668
Configuring VxSAS	673
Resetting the driver signing options	675
Configuring the cluster	676
Configuring Web console	688
Configuring notification	689
Configuring cluster disk groups and volumes	692
Creating a cluster disk group	693
Creating volumes	695
Installing and configuring SQL Server 2005 on the first node	701
Installing Microsoft SQL Server	701

Setting the startup mode of the SQL Server 2005 services	705
Preparing to install SQL Server 2005 on the second node	706
Stopping the SQL Server 2005 Service	706
Deporting the cluster disk group	706
Importing the cluster disk group	707
Adding drive letters to mount the volumes	707
Renaming shared SQL Server 2005 files	708
Installing SQL Server 2005 on the second node	709
Installing SQL Server	709
Removing shared SQL Server files	713
Setting the internal name of the clustered instance	713
Configuring the VCS SQL Server service group	716
Creating the primary system zone	721
Creating a SQL Server user-defined database	722
Creating new volumes	722
Creating a new SQL Server database	722
Adding VMDg and MountV resources	723
Verifying the installation in the primary zone	724
Creating a parallel environment in the secondary zone	725
Adding the systems in the secondary zone to the cluster	726
Setting up the Replicated Data Sets (RDS)	734
Configuring a hybrid RVG service group for replication	745
Creating the RVG service group	745
Configuring the RVG service group for RDC replication	746
Configuring the RVG Primary resources	756
Configuring the primary system zone for the RVG	759
Setting a dependency between the service groups	760
Adding the nodes from the secondary zone to the RDC	761
Adding the nodes from the secondary zone to the RVG service group	761
Configuring secondary zone nodes in the RVG service group	764
Configuring the IP resources for failover	764
Adding the nodes from the secondary zone to the SQL Server service group	766
Configuring the zones in the SQL Server service group	767
Verifying the RDC configuration	768
Bringing the service group online	768
Switching online nodes	768
Additional instructions for GCO disaster recovery	769

Section 5 Disaster Recovery

Chapter 15 Disaster recovery for SQL: Overview

What is a disaster recovery solution?	773
Why implement a disaster recovery solution?	774
Understanding replication	774
What needs to be protected in a SQL Server environment?	775
Typical disaster recovery configuration	775

Chapter 16 Deploying disaster recovery: New SQL Server 2000 installation

Tasks for a new disaster recovery installation of Microsoft SQL Server 2000	778
Reviewing the requirements	781
Disk space requirements	782
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	782
Reviewing the configuration	785
Supported disaster recovery configurations for service group dependencies	788
Configuring the storage hardware and network	789
Setting up the secondary site: Configuring SFW HA and setting up a cluster	792
Installing SFW HA	792
Setting Windows driver signing options	792
Installing Storage Foundation HA for Windows	794
Resetting the driver signing options	798
Configuring the cluster	799
Configuring Web console	811
Configuring notification	812
Verifying your primary site configuration	815
Setting up your replication environment	815
Setting up security for VVR	816
Requirements for EMC SRDF array-based hardware replication	819
Requirements for Hitachi TrueCopy array-based hardware replication	821
Assigning user privileges (secure clusters only)	823
Configuring disaster recovery with the DR wizard	824
Cloning the storage on the secondary site using the DR wizard (VVR replication option)	829

Creating temporary storage on the secondary site using the DR wizard (array-based replication)	833
Installing and configuring SQL Server 2000 on the first node (secondary site)	837
Installing Microsoft SQL Server	837
Setting SQL Server 2000 services to manual start	840
Preparing to install SQL Server on the second node (secondary site)	840
Stopping the SQL Server 2000 Service	841
Deporting the cluster disk group	841
Importing the cluster disk group	842
Mounting the volumes	842
Renaming shared SQL Server 2000 files	844
Installing SQL Server 2000 on the second node (secondary site)	844
Installing SQL Server	844
Removing shared SQL Server files	847
Setting the internal name of the clustered instance	848
Cloning the service group configuration from the primary to the secondary site	850
Configuring replication and global clustering	854
Configuring VVR replication and global clustering	854
Configuring EMC SRDF replication and global clustering	862
Configuring Hitachi TrueCopy replication and global clustering	865
Configuring global clustering only	868
Verifying the disaster recovery configuration	870
Establishing secure communication within the global cluster (optional)	872
Adding multiple DR sites (optional)	874
Recovery procedures for service group dependencies	875

Chapter 17

Deploying disaster recovery: New SQL Server 2005 installation

Tasks for a new disaster recovery installation of Microsoft SQL Server 2005	880
Reviewing the requirements	883
Disk space requirements	884
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	884
Reviewing the configuration	888
Supported disaster recovery configurations for service group dependencies	892

Configuring the storage hardware and network	892
Setting up the secondary site: Configuring SFW HA and setting up a cluster	895
Installing SFW HA	895
Setting Windows driver signing options	895
Installing Storage Foundation HA for Windows	897
Resetting the driver signing options	901
Configuring the cluster	902
Configuring Web console	914
Configuring notification	915
Verifying your primary site configuration	918
Setting up your replication environment	918
Setting up security for VVR	919
Requirements for EMC SRDF array-based hardware replication	922
Requirements for Hitachi TrueCopy array-based hardware replication	924
Assigning user privileges (secure clusters only)	926
Configuring disaster recovery with the DR wizard	927
Cloning the storage on the secondary site using the DR wizard (VVR replication option)	931
Creating temporary storage on the secondary site using the DR wizard (array-based replication)	935
Installing and configuring SQL Server 2005 on the first node (secondary site)	939
Installing SQL Server 2005 on the first node	939
Setting SQL Server 2005 services to manual start	943
Preparing to install SQL Server 2005 on the second node (secondary site)	944
Stopping the SQL Server 2005 Service	944
Deporting the cluster disk group	945
Importing the cluster disk group	945
Mounting the volumes	946
Renaming shared SQL Server 2005 files	947
Installing SQL Server 2005 on the second node (secondary site)	947
Installing SQL Server on the second node	947
Removing shared SQL Server files	951
Setting the internal name of the clustered instance	951
Cloning the service group configuration from the primary to the secondary site	954
Configuring replication and global clustering	958
Configuring VVR replication and global clustering	958
Configuring EMC SRDF replication and global clustering	966

	Configuring Hitachi TrueCopy replication and global clustering	969
	Configuring global clustering only	972
	Verifying the disaster recovery configuration	974
	Establishing secure communication within the global cluster (optional)	976
	Adding multiple DR sites (optional)	978
	Recovery procedures for service group dependencies	979
Chapter 18	Testing fault readiness by running a fire drill	
	About disaster recovery fire drills	983
	About the Fire Drill Wizard	984
	About post-fire drill scripts	985
	Tasks for configuring and running fire drills	986
	Prerequisites for a fire drill	986
	Fire Drill Wizard actions	988
	Preparing the fire drill configuration	990
	Running a fire drill	993
	Recreating a fire drill configuration that has changed	995
	Restoring the fire drill system to a prepared state	996
	Deleting the fire drill configuration	998
Section 6	Appendices	
Appendix A	Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation	
	Reviewing the prerequisites	1005
	Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	1005
	Reviewing the configuration	1009
	Sample configuration	1011
	Configuring the storage hardware and network	1012
	Installing Veritas Storage Foundation HA for Windows	1015
	Setting Windows driver signing options	1015
	Installing Storage Foundation HA for Windows	1017
	Configuring VxSAS	1021
	Resetting the driver signing options	1024
	Configuring the cluster	1024
	Configuring cluster disk groups and volumes	1041
	Creating a cluster disk group	1042
	Creating volumes	1044
	Installing and configuring SQL Server 2000 on the first node	1049

Installing Microsoft SQL Server	1049
Setting SQL Server 2000 services to manual start	1052
Preparing to install SQL Server on the second node	1052
Stopping the SQL Server 2000 Service	1053
Deporting the cluster disk group	1053
Importing the cluster disk group	1054
Adding drive letters to mount the volumes	1054
Renaming shared SQL Server 2000 files	1056
Installing SQL Server 2000 on the second node	1056
Installing SQL Server	1056
Removing shared SQL Server files	1059
Setting the internal name of the clustered instance	1060
Configuring the VCS SQL Server service group	1062
Creating a SQL Server user-defined database	1067
Creating new volumes	1067
Creating a new SQL Server database	1067
Adding VMDg and MountV resources	1068
Verifying the cluster configuration	1069
Creating a parallel environment on the secondary site	1070
Installing DR components on primary and secondary sites	1071

Appendix B Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation

Reviewing the prerequisites	1077
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	1077
Reviewing the configuration	1082
Sample configuration	1084
Configuring the storage hardware and network	1085
Installing Veritas Storage Foundation HA for Windows	1088
Setting Windows driver signing options	1088
Installing Storage Foundation HA for Windows	1090
Configuring VxSAS	1094
Resetting the driver signing options	1097
Configuring the cluster	1097
Configuring cluster disk groups and volumes	1114
Creating a cluster disk group	1115
Creating volumes	1117
Installing and configuring SQL Server 2005 on the first node	1122
Installing SQL Server 2005 on the first node	1122
Setting SQL Server 2005 services to manual start	1125
Preparing to install SQL Server 2005 on the second node	1126

Stopping the SQL Server 2005 Service	1126
Deporting the cluster disk group	1127
Importing the cluster disk group	1127
Adding drive letters to mount the volumes	1128
Renaming shared SQL Server 2005 files	1129
Installing SQL Server 2005 on the second node	1129
Installing SQL Server on the second node	1129
Removing shared SQL Server files	1133
Setting the internal name of the clustered instance	1133
Configuring the VCS SQL Server service group	1136
Creating a SQL Server user-defined database	1141
Creating new volumes	1141
Creating a new SQL Server database	1141
Adding VMDg and MountV resources	1142
Verifying the cluster configuration	1143
Creating a parallel environment on the secondary site	1144
Installing DR components on the primary and secondary sites	1145

Appendix C

Configuring the DR components (VVR and GCO) without using the DR wizard

Reviewing the prerequisites	1149
Setting up the Replicated Data Sets (RDS)	1149
Creating the VVR RVG service group	1160
Configuring the global cluster option for wide-area failover	1162
Prerequisites	1162
Linking clusters: Adding a remote cluster to a local cluster	1163
Converting a local service group to a global service group	1164
Bringing a global service group online	1166
Establishing secure communication within the global cluster (optional)	1167

Appendix D

Configuring an MSDTC service group for disaster recovery

Tasks for configuring MSDTC for disaster recovery	1169
Reviewing the prerequisites	1171
Reviewing the configuration	1171
Configuring cluster disk groups and volumes	1175
Creating a cluster disk group	1176
Creating volumes	1177
Mounting drives used by the MSDTC service group	1179
Creating an MSDTC service group	1180

Creating an MSDTC client	1182
Setting up the secondary site: Creating a parallel environment	1183
Installing DR components on the primary and secondary sites	1183

Index	1185
-------	------

Introduction

This section contains the following chapters:

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

Introducing Veritas Storage Foundation and High Availability Solutions for Microsoft SQL

This chapter includes the following topics:

- [“About the solutions guides”](#) on page 28
- [“About high availability”](#) on page 28
- [“About campus clusters”](#) on page 29
- [“About replicated data clusters”](#) on page 29
- [“About disaster recovery”](#) on page 29
- [“How this guide is organized”](#) on page 30

About the solutions guides

The *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* covers the following solutions for SFW HA with Microsoft SQL Server:

- High availability (HA)
- Campus clusters
- Replicated data clusters
- Disaster recovery (DR)

Solutions for Quick Recovery and Microsoft clustering solutions are in *Veritas Storage Foundation and High Availability Solutions Quick Recovery and Microsoft Clustering Solutions Guide for Microsoft SQL*.

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

About high availability

The term high availability (HA) refers to a state where data and applications are highly available because software or hardware is in place to maintain the continued functioning in the event of computer failure. High availability can refer to any software or hardware that provides fault tolerance, but generally the term has become associated with clustering.

Local clustering provides high availability through database and application failover. Veritas Storage Foundation HA for Windows (SFW HA) includes Veritas Storage Foundation and Veritas Cluster Server and provides the capability for local clustering.

Information about high availability for SQL Server includes procedures for installing and configuring clustered Microsoft SQL Server environments using SFW HA.

About campus clusters

Campus clusters are clusters 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.

Campus 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.

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 replicated data clusters

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage, to assure data access to all the nodes in a cluster.

The Replicated Data Cluster configuration provides both local high availability and disaster recovery functionality in a single VCS cluster. You can set up RDC in a VCS environment using Veritas Volume Replicator (VVR.)

An RDC exists within a single VCS cluster with a primary zone and a secondary zone, which can stretch over two buildings or data centers connected with Ethernet. In an RDC configuration, if an application or a system fails, the application is failed over to another system within the current primary zone. If the entire primary zone fails, the application is migrated to a system in the secondary zone (which then becomes the new primary).

About disaster recovery

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

Information about the disaster recovery solution for SQL Server includes procedures for installing, configuring, and testing clustered and replicated Microsoft SQL Server environments for disaster recovery using SFW HA.

How this guide is organized

The *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL* is organized to follow the workflow in the Solutions Configuration Center.

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

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

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

Using the Solutions Configuration Center

This chapter covers the following topics:

- [“About the Solutions Configuration Center”](#) on page 31
- [“Starting the Configuration Center”](#) on page 32
- [“Available options from the Configuration Center”](#) on page 32
- [“About running the Configuration Center wizards”](#) on page 39
- [“Following the workflow in the Configuration Center”](#) on page 40
- [“Solutions wizard logs”](#) on page 43

About the Solutions Configuration Center

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

- Microsoft Exchange Server 2003 and 2007
- Microsoft SQL Server 2000 and 2005
- Enterprise Vault Server (high availability solution only)
- Additional applications

You can use the Configuration Center and its wizards to set up your environment for any combination of the following solutions:

- 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 Veritas Cluster Server (SFW HA)
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data (available for Microsoft Exchange 2003 and 2007 and for Microsoft SQL Server 2005)
- Fire drill to test the fault readiness of a disaster recovery environment that uses VVR replication

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

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

Starting the Configuration Center

You can start the Configuration Center in two ways:

- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > Run**, type **scc** and click **OK**.

Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the Solution Guides listed in the right pane match the selected application.

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

[Figure 2-1](#) shows the choices available on a 32-bit system when you click Solutions for Microsoft Exchange.

Figure 2-1 Solutions Configuration Center for Microsoft Exchange

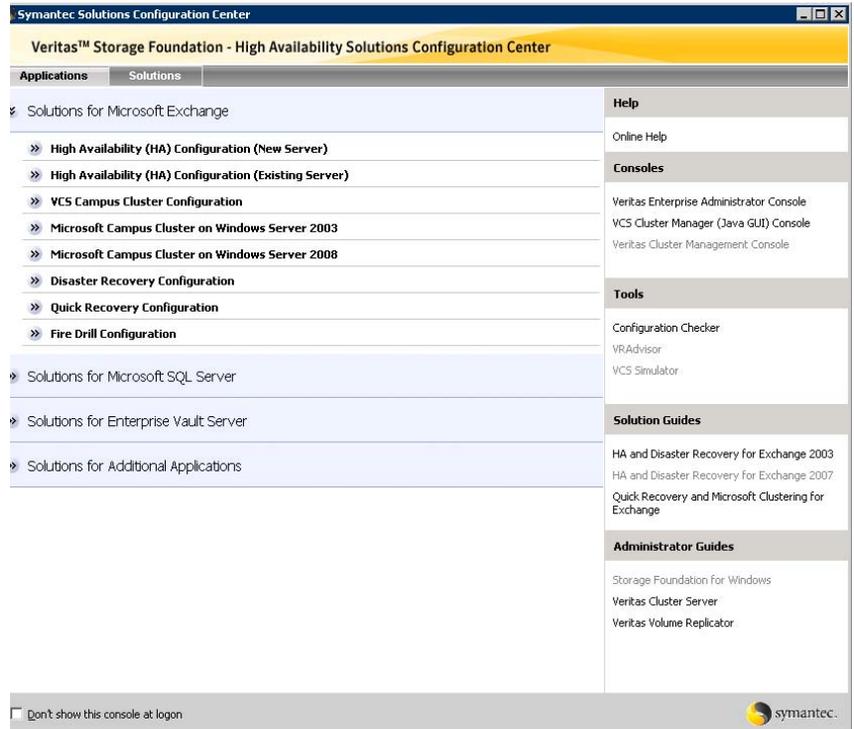


Figure 2-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 2-2 Solutions Configuration Center for Microsoft SQL Server

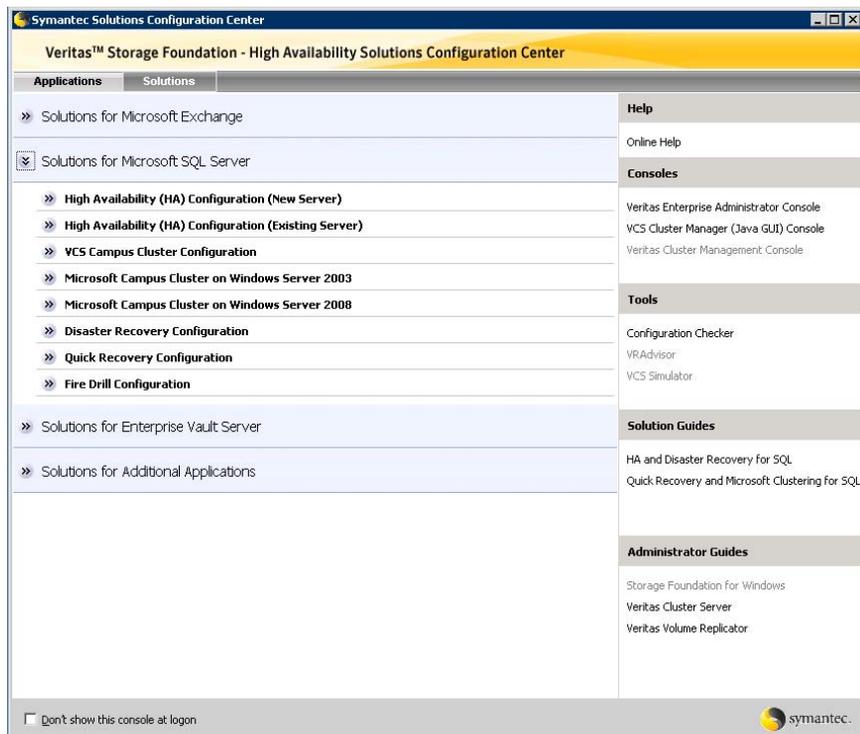


Figure 2-3 shows the choices available when you click Solutions for Enterprise Vault Server.

Figure 2-3 Solutions Configuration Center for Enterprise Vault Server

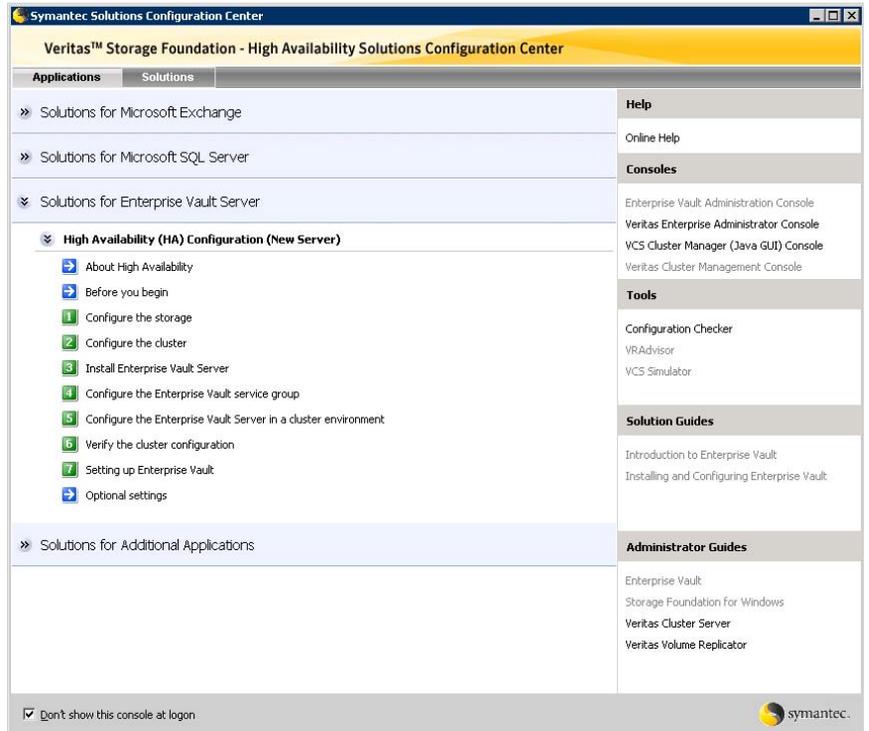
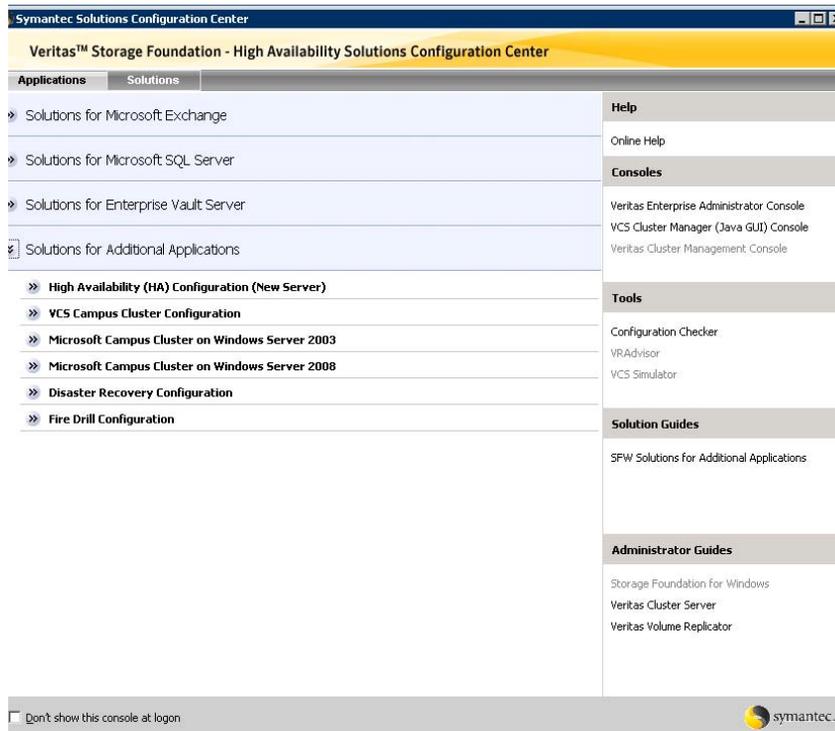


Figure 2-4 shows the choices available when you click Solutions for Additional Applications.

Figure 2-4 Solutions Configuration Center for additional applications



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

Figure 2-5 shows one of the steps for implementing high availability for Exchange.

Figure 2-5 Context-sensitive step for Exchange



Figure 2-6 shows one of the steps for implementing high availability for SQL Server.

Figure 2-6 Context-sensitive step for SQL Server

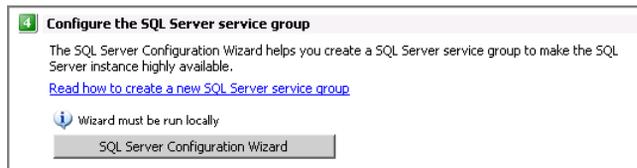


Figure 2-7 shows one of the steps for implementing high availability for Enterprise Vault Server.

Figure 2-7 Context-sensitive step for Enterprise Vault Server



Figure 2-8 shows one of the steps for implementing high availability for additional applications.

Figure 2-8 Context-sensitive step for additional applications

4 Configure the service group

Create a service group to make your application or server role highly available.

[Read how to create a File Share service group](#)

Wizard must be run locally

File Share Configuration Wizard

[Read how to create a Print Share service group](#)

Wizard must be run locally

Print Share Configuration Wizard

[Read how to create an IIS Server service group](#)

Wizard must be run locally

IIS Configuration Wizard

[Read how to create a Microsoft Virtual Server virtual machine service group](#)

Wizard must be run locally

MSVirtual Machine Configuration Wizard

[Read how to create an Oracle service group](#)

Wizard must be run locally

Oracle Agent Configuration Wizard

[Read how to create a service group for an application, process or service](#)

Wizard must be run locally

Application Configuration Wizard

About running the Configuration Center wizards

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

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

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

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

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
Exchange Setup Wizard	Installs and configures Exchange for the high availability environment If Exchange is already installed, refer to the documentation for further instructions.
Exchange Configuration Wizard	Configures the service group for Exchange high availability
SQL Server Configuration Wizard	Configures the service group for SQL Server high availability You must first install SQL Server on each node according to the instructions in the documentation.

Enterprise Vault Cluster Setup Wizard Configures the service group for Enterprise Vault Server high availability.

In addition, 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 FileShare for high availability.

Print Share Configuration Wizard Configures PrintShare for high availability.

IIS Configuration Wizard Configures IIS for high availability.

MSVirtual Machine Configuration Wizard Configures MS Virtual Machine for high availability.

Oracle Agent Configuration Wizard Configures Oracle for high availability.

Application Configuration Wizard Configures any other application service group for which application-specific wizards have not been provided.

Following the workflow in the Configuration Center

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

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

[Figure 2-9](#) shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

Figure 2-9 Workflow for configuring Exchange high availability

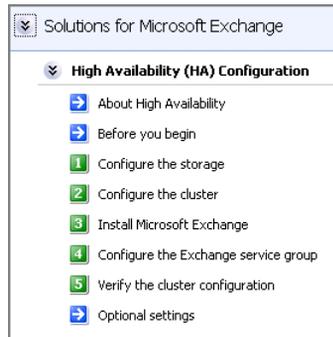


Figure 2-10 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

Figure 2-10 Workflow for configuring SQL Server high availability

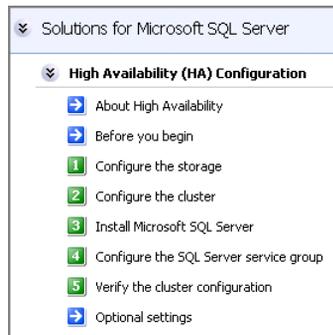


Figure 2-11 shows the high-level overview of the workflow steps for configuring high availability for Enterprise Vault Server from the Solutions Configuration Center.

Figure 2-11 Workflow for configuring high availability for Enterprise Vault Server

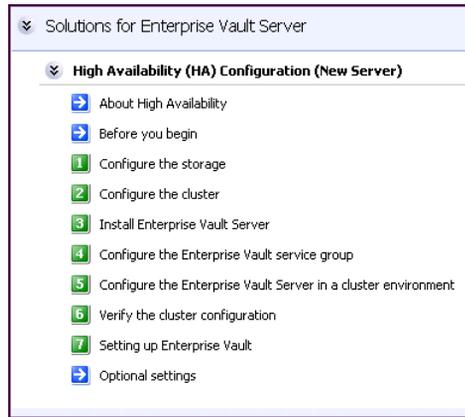
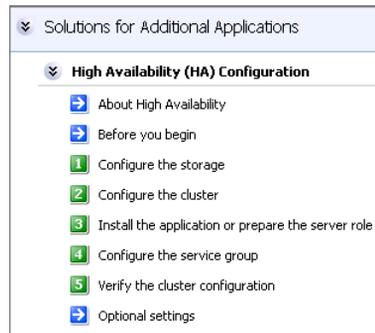


Figure 2-12 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

Figure 2-12 Workflow for configuring high availability for additional applications



Solutions wizard logs

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

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

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

Logs created by these three Solutions wizards are located in the following paths:

For Windows Server 2003:

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

For Windows Server 2008:

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


High Availability

This section includes the following chapters:

- [High availability for SQL: Overview](#)
- [Deploying SFW HA for high availability: New SQL Server 2000 installation](#)
- [Deploying SFW HA for high availability: New SQL Server 2005 installation](#)
- [Deploying SFW HA for high availability: Standalone SQL Servers](#)
- [Configuring an MSDTC service group for high availability](#)

High availability for SQL: Overview

This chapter includes the following topics:

- [“What is high availability?”](#) on page 47
- [“Why implement a high availability solution?”](#) on page 48
- [“How the agent makes SQL Server highly available”](#) on page 48

What is high availability?

The term *High Availability* 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. This section focuses on configurations that use Veritas Storage Foundation HA for Windows. Storage Foundation HA for Windows (SFW HA) includes Veritas Storage Foundation and Veritas Cluster Server.

Local clustering provides high availability (HA) through database and application failover. This solution provides local recovery in the event of application, operating system, or hardware failure, and minimizes planned and unplanned application downtime. The High Availability section includes procedures for installing and configuring clustered Microsoft SQL Server environments using Veritas Storage Foundation HA for Windows.

Setting up the clustered environment is also the first step in creating a disaster recovery solution. Some installation and configuration options in this section are identified as required for disaster recovery only. These options apply only if you intend to set up a secondary site for wide-area disaster recovery.

A cluster is a group of independent computers working together to ensure that mission-critical applications and resources are as highly available as possible. The group 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.

Why implement a high availability solution?

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

Using Veritas Storage Foundation HA for Windows as a local high availability solution paves the way for a wide-area disaster recovery solution in the future. A high availability solution is built on top of a backup strategy and provides the following benefits:

- Reduces planned and unplanned downtime.
- Serves as a local and wide-area failover (rather than load-balancing) solution. Enables failover between sites or between clusters.
- Manages applications and provides an orderly way to bring processes online and take them offline.
- Consolidates hardware in larger clusters. The HA environment accommodates flexible fail over policies, active-active configurations, and shared standby servers for SQL Server.

How the agent makes SQL Server highly available

The VCS enterprise agent for SQL Server detects an application failure if a configured virtual server becomes unavailable. When this occurs, the SQL Server service group is failed over to the next available system in the service group's system list. The configured SQL services and virtual server are started on the new system.

Deploying SFW HA for high availability: New SQL Server 2000 installation

This chapter covers the following topics:

- [“Tasks for a new HA installation of SQL Server 2000”](#) on page 50
- [“Reviewing the requirements”](#) on page 55
- [“Reviewing the configuration”](#) on page 59
- [“Reviewing considerations for Active-Active configurations”](#) on page 67
- [“Configuring the storage hardware and network”](#) on page 70
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 73
- [“Configuring cluster disk groups and volumes for SQL Server 2000”](#) on page 79
- [“Configuring the cluster”](#) on page 88
- [“About installing multiple instances”](#) on page 104
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 105
- [“Preparing to install SQL Server 2000 on the second node”](#) on page 108
- [“Installing and configuring SQL Server 2000 on the second node”](#) on page 112
- [“Creating a SQL Server user-defined database”](#) on page 118
- [“Completing configuration steps in SQL Server”](#) on page 119

- [“Configuring the VCS SQL Server 2000 service group”](#) on page 121
- [“Verifying the SQL Server 2000 cluster configuration”](#) on page 127
- [“Determining additional steps needed”](#) on page 128
- [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 133

Tasks for a new HA installation of SQL Server 2000

You can install and configure a new Veritas Storage Foundation HA environment for SQL Server 2000 in the following ways:

Active-Passive	<p>One SQL instance per node with one to one failover capabilities.</p> <p>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.</p>
Active-Active	<p>Multiple SQL instances per cluster node.</p> <p>For example, in a two-node cluster with two SQL 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.</p>

Note: Some installation and configuration options in this section are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2000.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for high availability”](#) on page 355.

If you have an existing SQL Server 2000 deployment on a standalone server, refer to [“Deploying SFW HA for high availability: Standalone SQL Servers”](#) on page 135.

Tasks for an Active-Passive configuration

Table 4-1 outlines the high-level objectives and the tasks to complete each objective for an active-passive configuration.

Table 4-1 SQL Server 2000: Active-Passive configuration tasks

Objective	Tasks
“Reviewing the requirements” on page 55	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 59	<ul style="list-style-type: none"> ■ Understanding active-passive configuration ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 70	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Installing Veritas Storage Foundation HA for Windows” on page 73	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2000” on page 79	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Configuring the cluster” on page 88	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2000 on the first node” on page 105	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Setting SQL Server services to manual start

Table 4-1 SQL Server 2000: Active-Passive configuration tasks (Continued)

Objective	Tasks
“Preparing to install SQL Server 2000 on the second node” on page 108	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Removing shared SQL files
“Installing and configuring SQL Server 2000 on the second node” on page 112	Installing and configuring SQL Server 2000
“Setting the internal name of the clustered instance” on page 116	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 118	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 119	Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2000 service group” on page 121	Creating a SQL Server service group using the SQL Server Configuration Wizard
“Verifying the SQL Server 2000 cluster configuration” on page 127	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes

Tasks for an Active-Active configuration

Table 4-2 outlines the high-level objectives and the tasks to complete each objective for an active-active configuration.

Table 4-2 SQL Server 2000: Active-Active configuration tasks

Objective	Tasks
“ Reviewing the requirements ” on page 55	Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 59	<ul style="list-style-type: none"> ■ Understanding active-active configuration ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 70	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Installing Veritas Storage Foundation HA for Windows ” on page 73	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“ Configuring cluster disk groups and volumes for SQL Server 2000 ” on page 79	<ul style="list-style-type: none"> ■ Creating a dynamic disk group for each instance ■ Creating dynamic volumes for the SQL system database, user databases, transaction logs, and replicated registry keys for each instance
“ Configuring the cluster ” on page 88	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“ Installing and configuring SQL Server 2000 on the first node ” on page 105	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Setting SQL Server services to manual start

Table 4-2 SQL Server 2000: Active-Active configuration tasks (Continued)

Objective	Tasks
“Preparing to install SQL Server 2000 on the second node” on page 108	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2000 on the second node” on page 112	Installing SQL Server on the second node
“Setting the internal name of the clustered instance” on page 116	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 118	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 119	Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2000 service group” on page 121	<ul style="list-style-type: none"> ■ Creating a SQL Server service group using the SQL Server Configuration Wizard ■ Ensuring the priority order of the systems is set up in reverse order for each instance
“Verifying the SQL Server 2000 cluster configuration” on page 127	<ul style="list-style-type: none"> ■ Simulating fail over ■ Switching online nodes
“Determining additional steps needed” on page 128	Repeating the installation and configuration steps for the next instance

Reviewing the requirements

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation.

Disk space requirements

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

[Table 4-3](#) estimates disk space requirements for SFW HA.

Table 4-3 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 for 64-bit Itanium (IA64)■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 59.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW

HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

Review the information for the configurations you have planned:

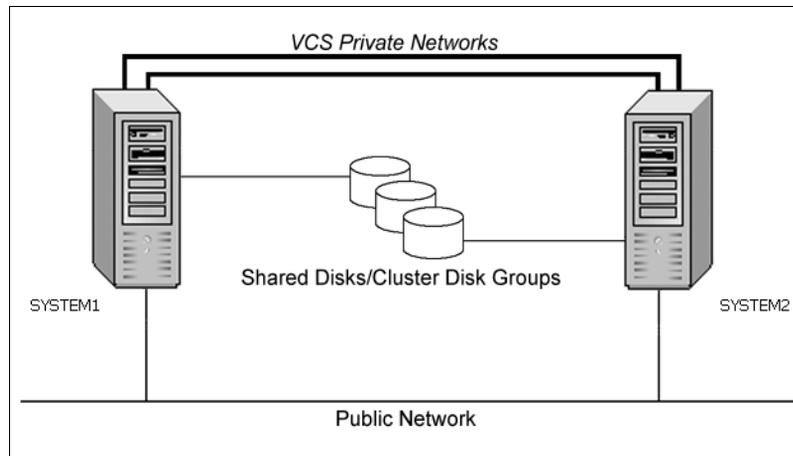
- [Active-Passive configuration](#)
- [Active-Active configuration](#)
- [Disaster recovery configuration](#)

Active-Passive configuration

In a typical example of a high availability cluster, you create a virtual server in an Active-Passive SQL configuration. 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.

Figure 4-1 illustrates a typical Active-Passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 4-1 Active-Passive configuration



Sample Active-Passive configuration

A sample setup is used to illustrate the installation and configuration tasks.

The following names describe the objects created and used during the installation and configuration:

Table 4-4 Active-Passive configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files

Object Name	Description
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	virtual SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
INST1_SG	SQL Server service group

IP addresses for sample Active-Passive configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there is one SQL Server virtual server. Therefore you would need one virtual server IP address. If you want to use the VCS Web Console or the notification service, you require a cluster IP address (the cluster IP address is also used by the Global Cluster Option for disaster recovery).

See “Network requirements” in [“Reviewing the requirements”](#) on page 55.

See [“Disaster recovery configuration”](#) on page 64.

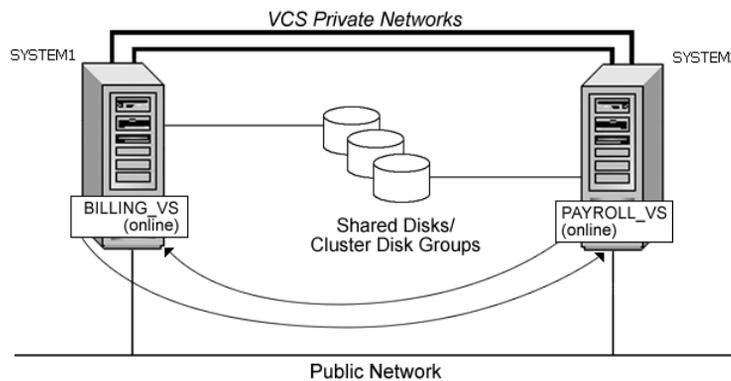
Active-Active configuration

In an Active-Active SQL Server configuration, several instances are intended to run on a single node when necessary. A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

SQL Server 2000 supports up to 16 independent instances of SQL Server 2000 to run on a single machine. A SQL Server 2000 instance can fail over to any of the other nodes on its system list. You can choose an Active-Active SQL Server 2000 configuration where several instances are intended to run on a single node. However, remember that you must configure failover nodes so that a single node can never host more than 16 instances.

The following figure illustrates a two node Active-Active configuration. The SQL Server databases are configured on the shared storage on volumes contained in cluster disk groups. Each SQL Server virtual server is configured in a separate SQL Server service group. Each service group can fail over to the other node in the cluster.

Figure 4-2 Active-Active configuration



For example, consider a two-node cluster hosting two SQL Server virtual servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to

SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

SQL Virtual Server	Service Group	System List
BILLING_VS	BILLING_SG	SYSTEM1, SYSTEM2
PAYROLL_VS	PAYROLL_SG	SYSTEM2, SYSTEM1

Sample Active-Active configuration

A sample setup is used to illustrate the installation and configuration tasks for two instances of SQL Server, Billing and Payroll. During normal operation, one instance will be online on each of the two servers. If a failure occurs, the instance on the failing node will be brought online on the other server, resulting in two instances running on one server.

The following names describe the objects created and used during the installation and configuration:

Table 4-5 Active-Active configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	server names
BILLING_DG	cluster disk group for the billing instance
PAYROLL_DG	cluster disk group for the payroll instance
BILLING_VS_SYS_FILES	volume for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume for the SQL Server system data files for the payroll instance
BILLING_DATA	volume for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume for a SQL Server user-defined database log file for the billing instance
PAYROLL_LOG	volume for a SQL Server user-defined database log file for the payroll instance

Object Name	Description
BILLING_REGREP	volume for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume for the list of registry keys replicated among the nodes for the payroll instance
SQL_CLUS1	virtual SQL Server cluster
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL Server name for the billing instance
PAYROLL_VS	virtual SQL Server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

IP addresses for sample Active-Active configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there are two virtual servers: BILLING-VS and PAYROLL-VS. Therefore, you would need two virtual server IP addresses. If you want to use the VCS Web Console or the notification service, you require a cluster IP address (the cluster IP address is also used by the Global Cluster Option for disaster recovery).

See “Network requirements” under “[Reviewing the requirements](#)” on page 55.

See “[Disaster recovery configuration](#)” on page 64.

Disaster recovery configuration

You may be preparing to configure both a primary site and a secondary site for disaster recovery.

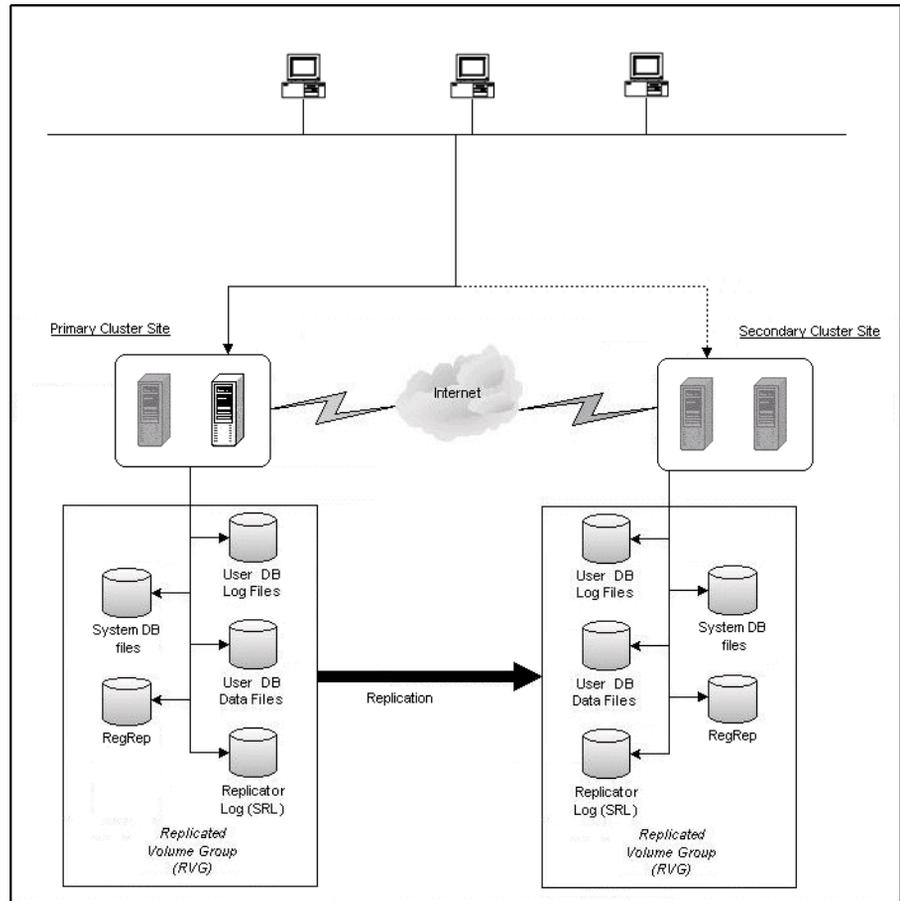
[Figure 4-3](#) illustrates a typical Active-Passive disaster recovery configuration using Veritas Volume Replicator (VVR).

In the example illustration, the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting

up the Replicated Volume Group (RVG). The SQL Server application data is stored on the volumes that are under the control of the RVG.

If the Microsoft SQL Server server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Figure 4-3 Typical disaster recovery configuration with VVR



Sample disaster recovery configuration

The sample disaster recovery setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

A sample setup is used to illustrate the installation and configuration tasks:

Table 4-6 Disaster Recover configuration objects

Object Name	Description
Primary Site	
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
Secondary Site	
SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.
DR Components	
INST1_REPLOG	replicator log volume required by VVR
INST1_DB1_RDS	RDS (Replicated Data Set) name
INST1_DB1_RVG	RVG (Replicated Volume Group) name
INST1_DB1_RVG_SG	Replication service group

IP addresses for sample disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration.

See “Network requirements” under “[Reviewing the requirements](#)” on page 55.

You specify the following addresses during the replication process:

Table 4-7 IP addresses required for DR configuration

IP address	Description
SQL virtual server IP address	For a disaster recovery configuration, the virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.
Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.
Replication IP address	You need two IP addresses per application instance, one for the primary site and one for the secondary site.

Reviewing considerations for Active-Active configurations

For an Active-Active configuration or in other cases when you are installing multiple instances of SQL Server on the same system, there are special considerations.

To assist you in planning your deployment, the considerations are summarized in the following topics:

- [“Key information for Active-Active configurations”](#) on page 68
- [“Following the workflow in the Solutions Configuration Center”](#) on page 69

Key information for Active-Active configurations

The following table summarizes key information about Active-Active configurations and multiple instances and cross-references additional information:

Table 4-8 Key information for Active-Active configuration

IP address	Description
Configuring disk groups and volumes	<p>Create a separate set of cluster disk groups and volumes for each instance. You can create all the disk groups and volumes at one time or create them as a separate step for each instance.</p> <p>See “Considerations for disk groups and volumes for multiple instances” on page 81.</p>
Configuring the cluster	<p>If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. That way, you do not need to run the wizard again later to add the nodes.</p> <p>See “Configuring the cluster” on page 88.</p>
Installing and configuring SQL Server	<p>Assign a unique instance name, virtual server name, and port to each instance.</p> <p>“About installing multiple instances” on page 104.</p> <p>“Setting the internal name of the clustered instance” on page 116.</p> <p>“Assigning ports for multiple SQL Server instances” on page 120.</p>
Configuring the service group	<p>For an active/active configuration, create a separate service group for each instance. Each service group must have a unique name and virtual IP address. There are also special considerations for specifying the priority order of systems for failover.</p> <p>See “Service group requirements for active-active configurations” on page 121.</p>

Following the workflow in the Solutions Configuration Center

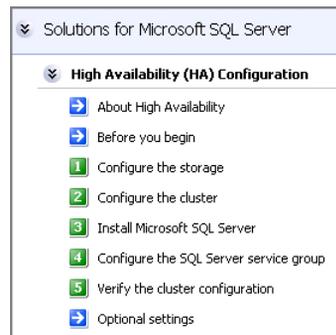
The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA environment for one or more instances of SQL Server 2000, in either an active-passive or active-active configuration.

Figure 4-4 shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

If you are setting up multiple instances of SQL in the cluster, you may find it helpful to use the Configuration Center as follows:

- Under High Availability (HA) Configuration, complete all the steps for the first instance.
- For the next instance:
 - For step 1, Configure the storage: If you configured disk groups and volumes for the instance earlier, verify that they are available and continue with step 2.
 - For step 2, Configure the cluster: If you configured the nodes as part of the cluster earlier, as recommended, continue with step 3 and complete all subsequent steps.

Figure 4-4 Configuration steps in the Solutions Configuration Center



See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 4-9](#) on page 73 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 4-9 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 74.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

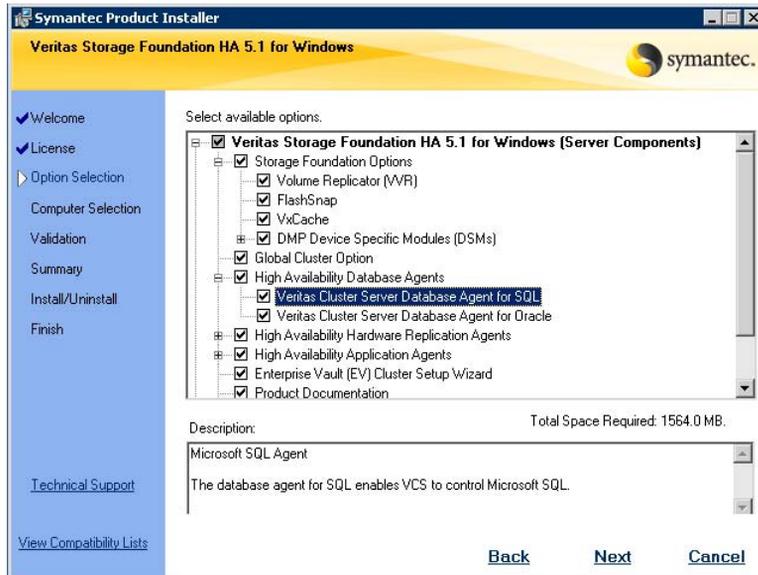
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

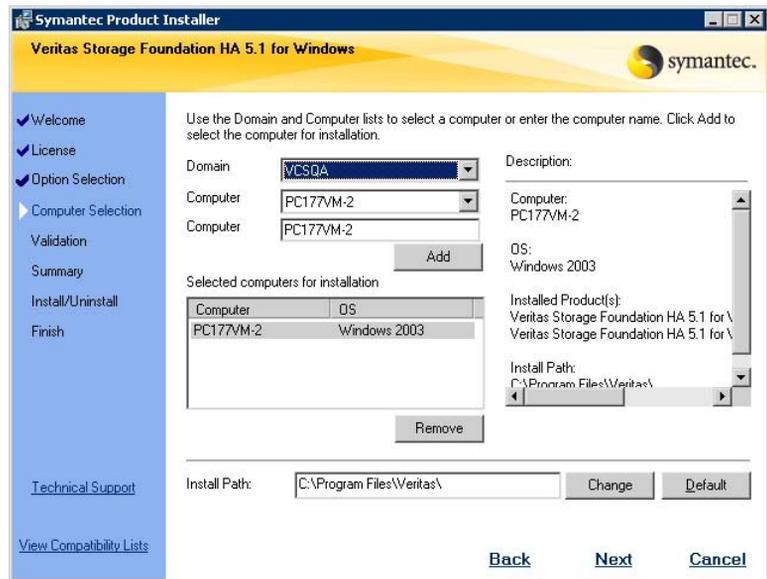
Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option Required for a disaster recovery configuration only.

Veritas Volume Replicator For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

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

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Configuring cluster disk groups and volumes for SQL Server 2000

You must create a cluster disk group and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [“About cluster disk groups and volumes”](#) on page 79
- [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 80
- [“Sample disk group and volume configuration”](#) on page 80
- [“Considerations for disk groups and volumes for multiple instances”](#) on page 81
- [“Considerations for volumes for a disaster recovery configuration”](#) on page 82
- [“Creating a cluster disk group”](#) on page 83
- [“Creating volumes”](#) on page 84

About cluster disk groups and volumes

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: Create a cluster disk group and volumes on only one node of a cluster. You make the volumes accessible 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node of the cluster where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

For additional disk group and volume requirements, see the following topics:

- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)

Considerations for disk groups and volumes for multiple instances

For an active-active configuration or other cases where you are setting up multiple SQL instances in the cluster, you create a separate set of cluster disk groups and volumes for each instance.

For example, if you have a Billing instance and a Payroll instance, you could create the following disk groups and volumes.

For the Billing instance, create the following:

- BILLING_DG: a cluster disk group for the volumes related to the Billing instance
- BILLING_DATA_FILES: volume for the SQL Server system data files
- BILLING_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Billing instance
- BILLING_DB1_VOL: volume for the user database files
- BILLING_DB1_LOG: volume for the user database log files

For the Payroll Instance, create the following:

- PAYROLL_DG: a cluster disk group for the volumes related to the Payroll instance
- PAYROLL_DATA_FILES: volume for the SQL Server system data files
- PAYROLL_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Payroll instance
- PAYROLL_DB1_VOL: volume for the user database files
- PAYROLL_DB1_LOG: volume for the user database log files

You can choose either of the following:

- Set up disk groups and volumes for all instances at one time.
- Set up disk groups and volumes for the current instance only and complete all configuration steps for this instance. Then return to this step for the next instance.

Considerations for volumes for a disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- A disaster recovery configuration with VVR requires a Storage Replicator Log (SRL) volume 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. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.
- Symantec recommends that for replication considerations, you create a separate volume for tempdb, for example, INST1_TEMPDB, within the system database disk group. When you later configure replication for disaster recovery, you replicate that disk group but exclude the tempdb volume from the replication.

It would waste bandwidth to replicate tempdb because the data is transitory and is not needed for DR site recovery.

You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume. See "[Moving the tempdb database if using VVR for disaster recovery](#)" on page 119.

- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

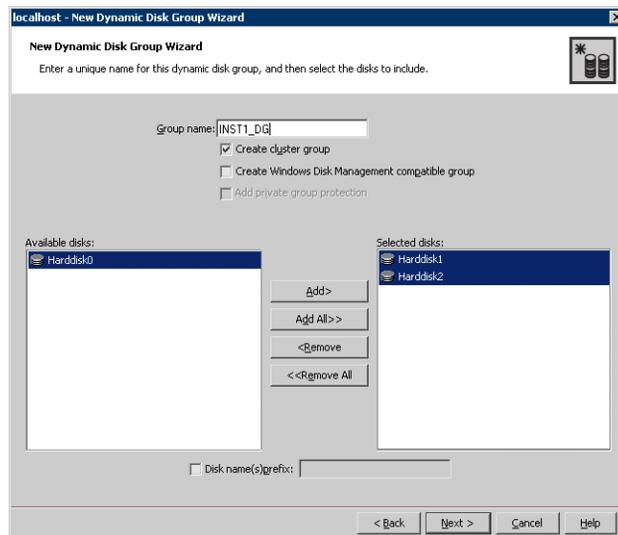
Creating a cluster disk group

Create a cluster disk group on the first node where the SQL instance is being installed. Repeat the procedure if you want to create additional disk groups.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

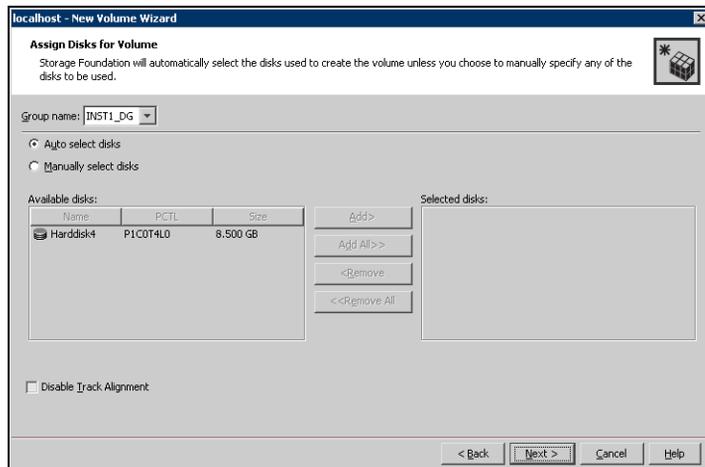
Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.

See “[Considerations for volumes for a disaster recovery configuration](#)” on page 82

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 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
 You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

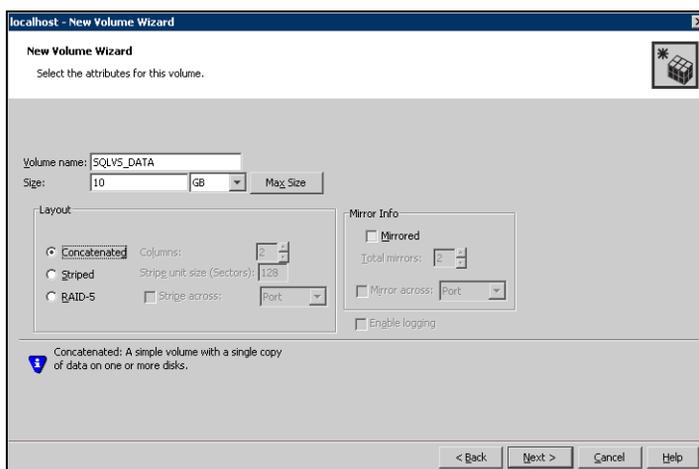


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove**

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

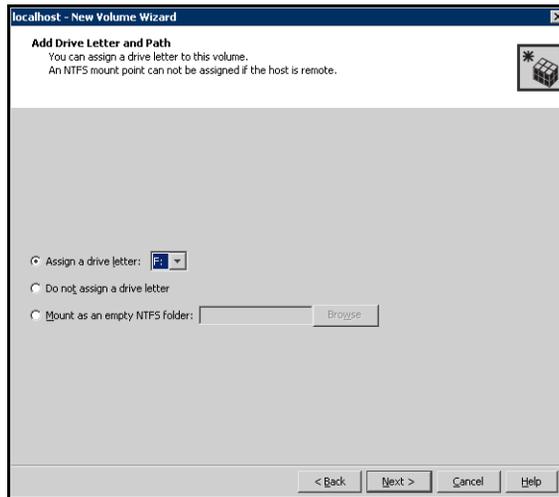
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

- 8 Click **Next**.
- 9 Specify the parameters of the volume.



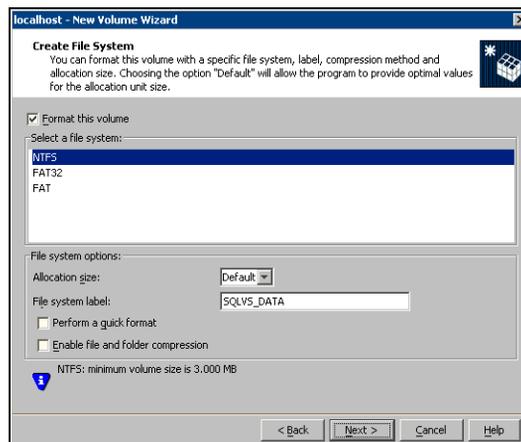
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 - To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the **Default**.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

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

14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses a static IP address and that name resolution is configured for each node.
- Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 55.

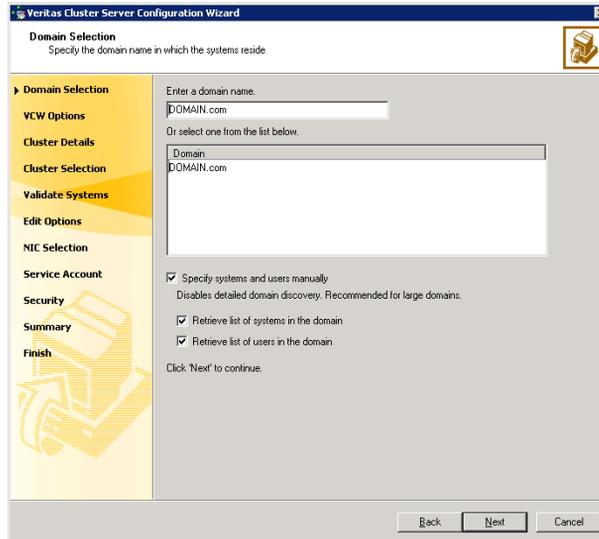
Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or removing or modifying cluster configurations.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



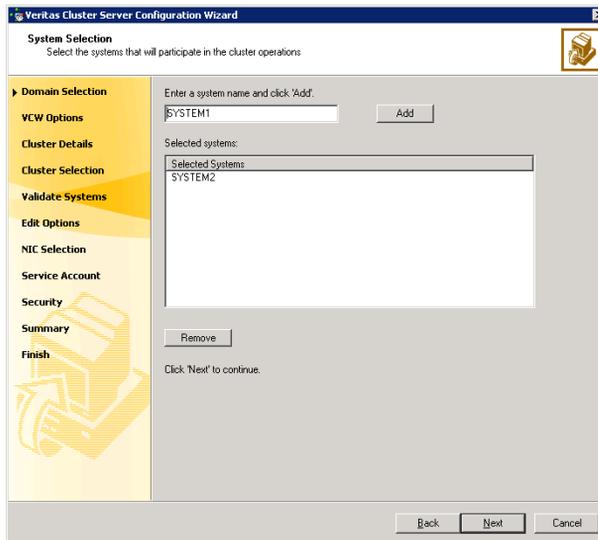
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
 Proceed to [step 8](#) on page 91.

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

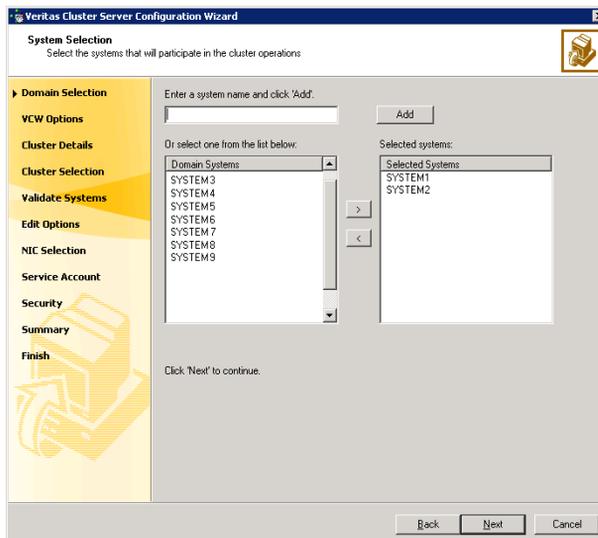
- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
 If you chose to retrieve the list of systems, proceed to [step 6](#) on page 90.
 Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 91.

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



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

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

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

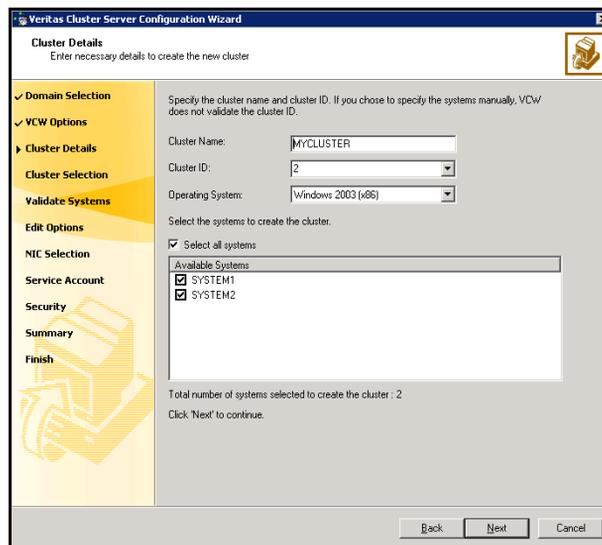
- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

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

Click **Next** to proceed.

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

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



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 89 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

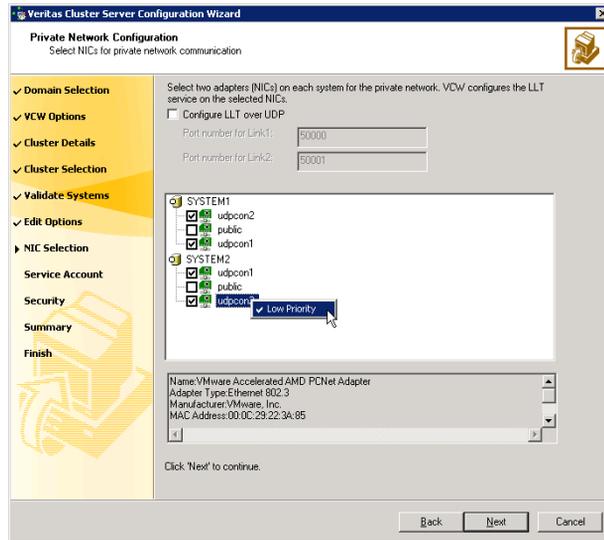
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 91, proceed to the next step. Otherwise, proceed to [step 12](#) on page 95.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

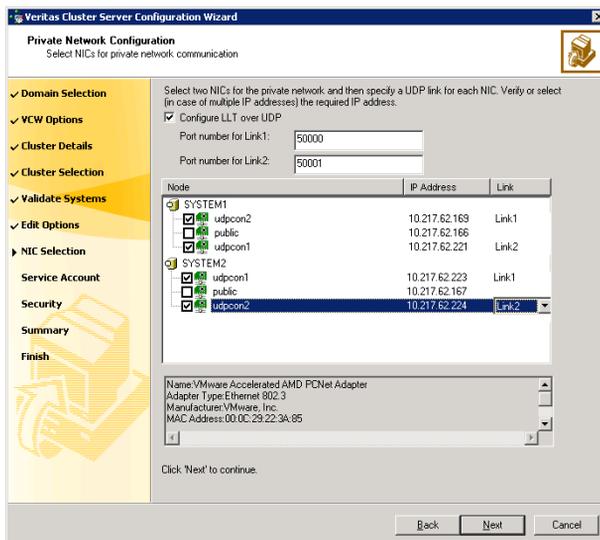
- To configure the VCS private network over Ethernet



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

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

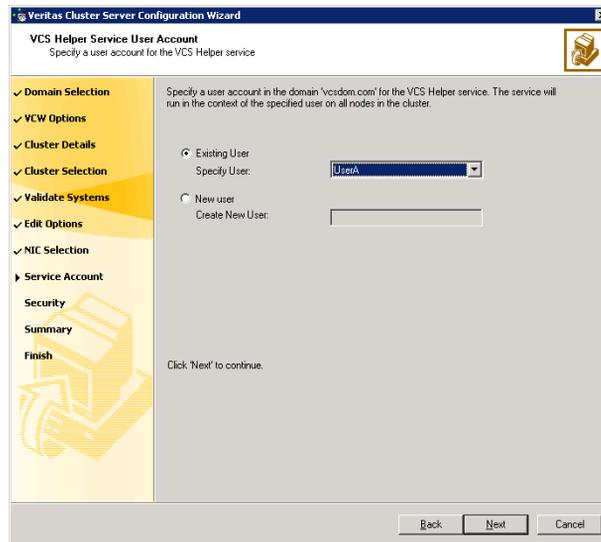
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



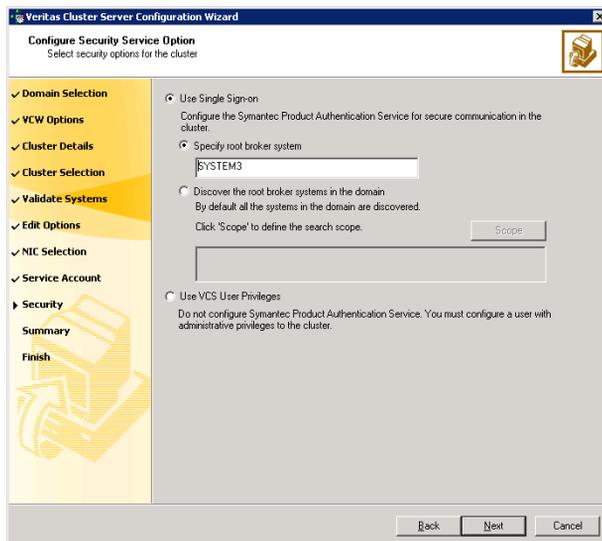
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 89, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
 Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

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

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 4-10](#) on page 97 contains some more examples of search criteria.

Table 4-10 Search criteria examples

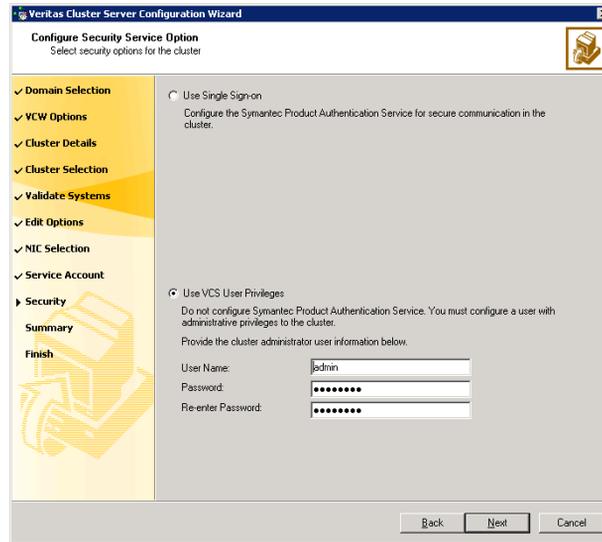
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.

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

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

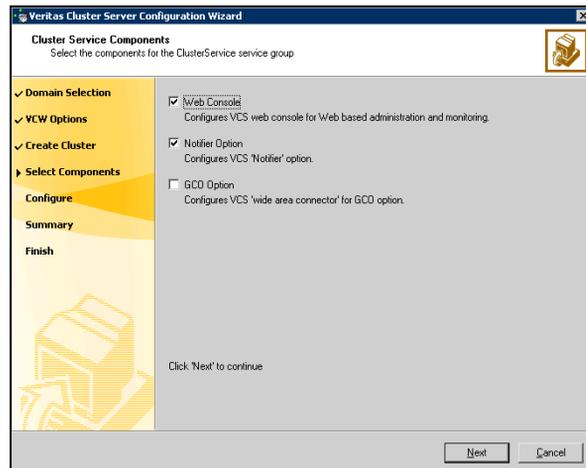
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 100.

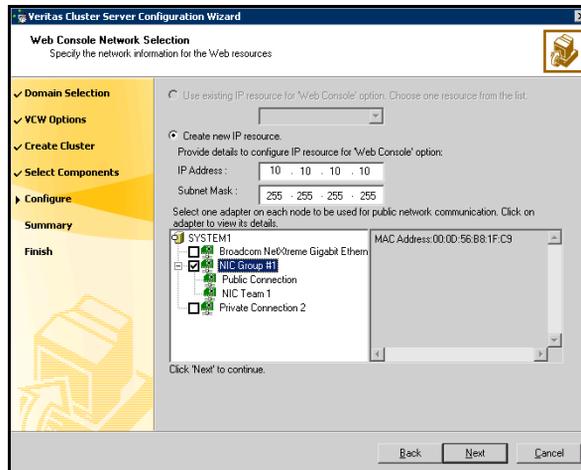
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See “[Configuring notification](#)” on page 101.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



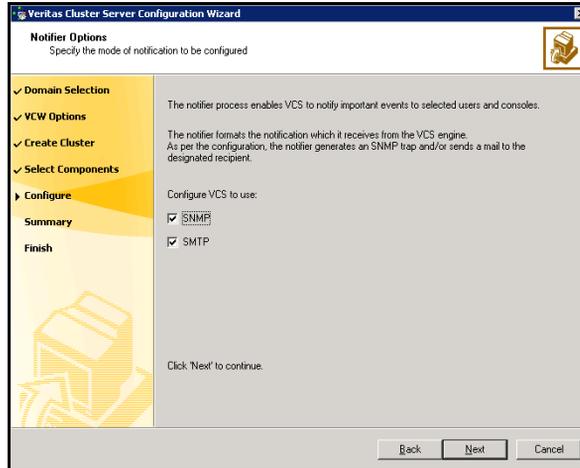
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 101.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

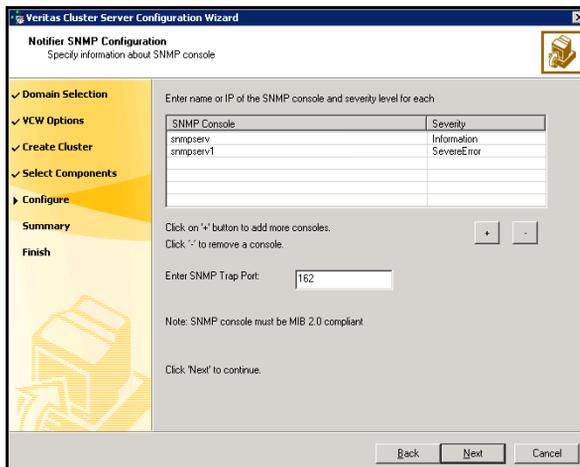
To configure notification

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

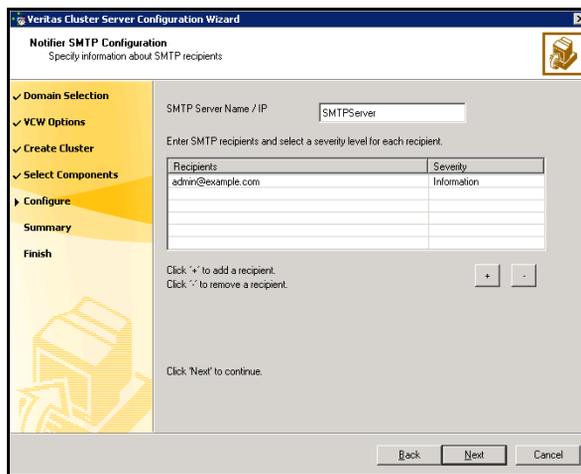


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

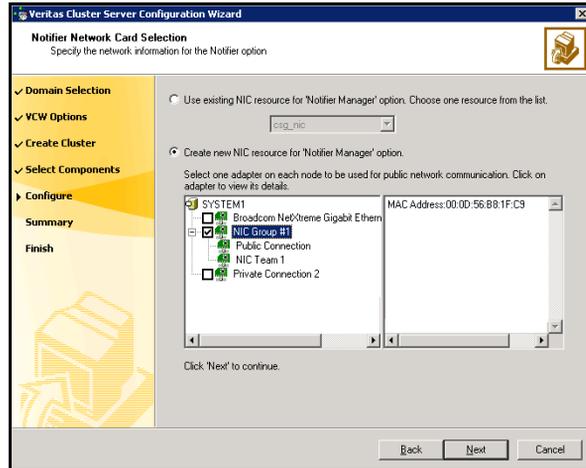


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

About installing multiple instances

If you are installing multiple instances of SQL Server on the same system, as in an Active-Active cluster configuration, some additional requirements apply. Procedures include these requirements where necessary. The following summary is provided for your review to assist you in planning the installation:

- Symantec recommends that you follow all steps for installing and setting up high availability for the first instance before you begin installing the next instance.
- Multiple instances of SQL Server must be installed in the same order on every node in the cluster. For example, if you install INST1 on SYSTEM1 and then on SYSTEM2, then install INST2 on SYSTEM1 and then on SYSTEM2.
- Assign a unique name to each instance. When installing SQL Server on additional nodes for the same instance, use the same instance name.
- Set a unique virtual server name as the internal name for each clustered instance.
- Assign a unique port number for each instance.

Installing and configuring SQL Server 2000 on the first node

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node for this SQL instance and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 110 and “[Adding drive letters to mount the volumes](#)” on page 110.

Complete the following procedures to install and configure this instance of Microsoft SQL Server 2000:

- [Installing Microsoft SQL Server 2000](#)
- [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server 2000

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.

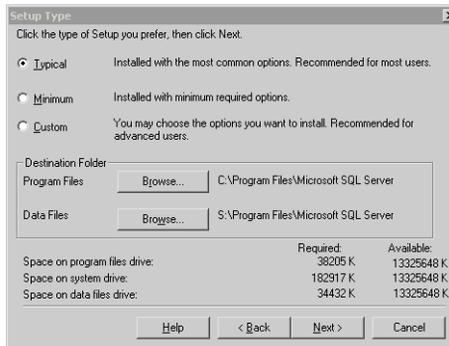
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

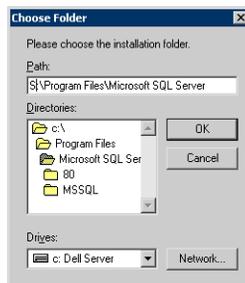
To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.

- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

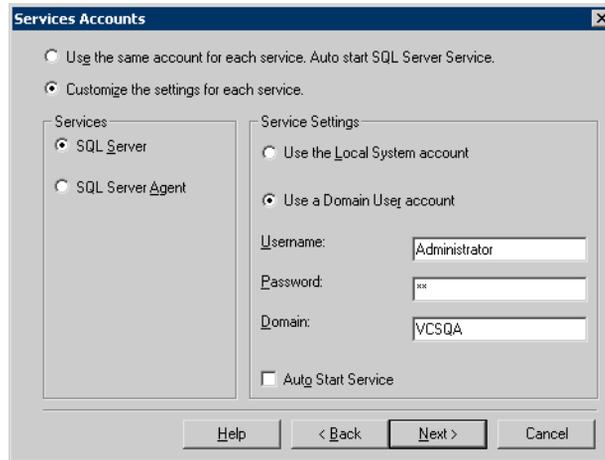


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

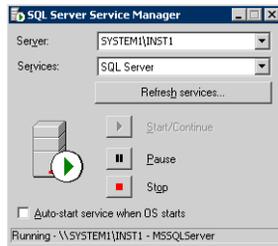
Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Preparing to install SQL Server 2000 on the second node

Follow the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:

- [“Stopping the SQL Server 2000 service”](#) on page 109
- [“Deporting the cluster disk group”](#) on page 109
- [“Importing the cluster disk group”](#) on page 110
- [“Adding drive letters to mount the volumes”](#) on page 110
- [“Renaming shared SQL Server 2000 files”](#) on page 112

Stopping the SQL Server 2000 service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.

- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

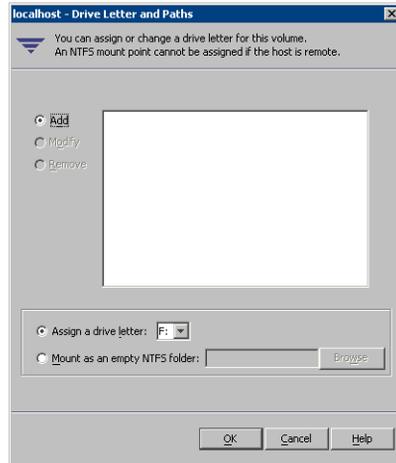
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2000 on the second node

Use the following procedures to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2000 on the second node”](#) on page 112
- [“Removing shared SQL Server 2000 files”](#) on page 115

Installing SQL Server 2000 on the second node

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 110 and [“Adding drive letters to mount the volumes”](#) on page 110.

Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

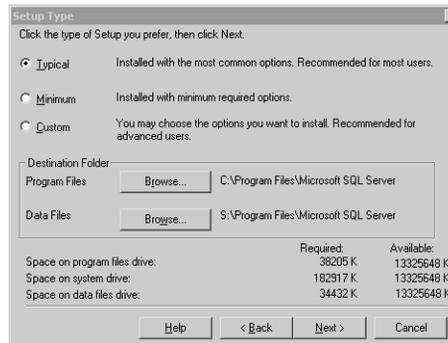
Multiple instances of SQL Server must be installed in the same order on every node of the cluster.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

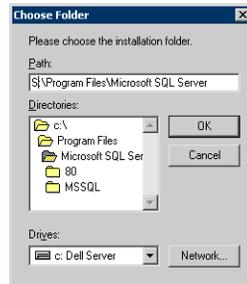
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

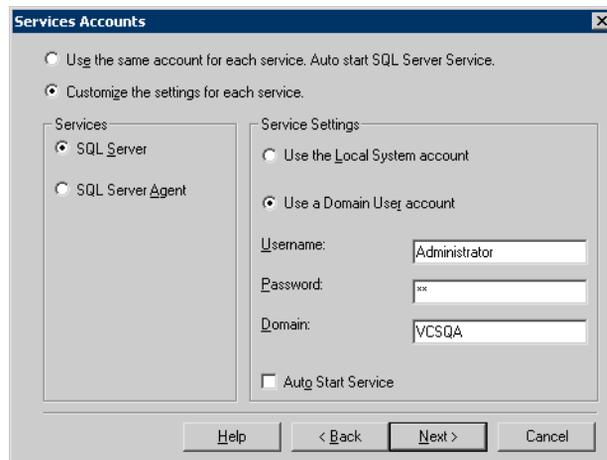
- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.

- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server 2000 on the second node](#)” on page 108 and “[Installing and configuring SQL Server 2000 on the second node](#)” on page 112 on any additional nodes.

Removing shared SQL Server 2000 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

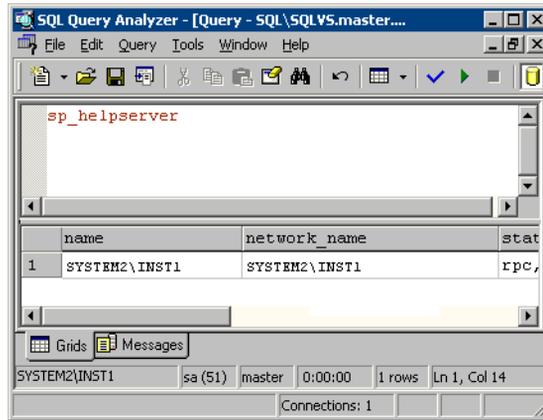
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`“sp_dropserver ‘System_Name\Instance_Name.’”`
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
`“sp_dropserver ‘SYSTEM2\INST1.’”`
 For example, for a default instance:
`“sp_dropserver ‘SYSTEM1.’”`
 - Press F5.
- 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter
`"sp_addserver 'Virtual_Server_Name\Instance_Name', local"`
For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.
 - Press F5.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database.

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See ["Creating volumes"](#) on page 84.

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.
- 8 Configure whatever other options are required for your database.
- 9 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See ["Completing configuration steps in SQL Server"](#) on page 119.
- 10 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See ["Modifying the SQL 2000 service group to add VMDg and MountV resources"](#) on page 133.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 119.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See “[Assigning ports for multiple SQL Server instances](#)” on page 120.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating volumes](#)” on page 84.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Assigning ports for multiple SQL Server instances

If you are running multiple SQL Server instances, you must assign a different port to each named instance. You can assign static or dynamic ports.

Refer to the Microsoft Knowledge Base for the instructions on assigning ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

If you wish to change the port after configuring the SQL service group, you must perform the steps in the following order:

- Bring the SQL service group online or partially online (up to the registry replication resource) on a cluster node.
- On the node on which the SQL service group is online or partially online, change the port assigned to the SQL instance. Refer to the instructions mentioned in the Microsoft Knowledge Base article specified earlier.
- Take the SQL service group offline on the node, and then bring it online again. The configuration changes will be replicated to the remaining cluster nodes.

Configuring the VCS SQL Server 2000 service group

A VCS SQL Server service group is used to bring a SQL Server 2000 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Read the following topics:

- [Service group requirements for active-active configurations](#)
- [Prerequisites for configuring the service group](#)
- [Creating the SQL Server 2000 service group](#)

Service group requirements for active-active configurations

Note the following requirements for active-active configurations:

- For an active-active configuration, you must create a separate service group for each instance.
- Each service group that you create must have a unique service group name and virtual IP address.
- For an active-active configuration, when you specify the priority order of systems, reverse the order for each service group so that the active system and failover system are opposite for each instance. For example, if you have two instances and two systems, you would set the priority order as follows:

INSTANCE 1	Priority order: SYSTEM1 SYSTEM2
INSTANCE2	Priority order: SYSTEM2 SYSTEM1

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 73.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 88.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group. See [“Importing the cluster disk group”](#) on page 110 and [“Adding drive letters to mount the volumes”](#) on page 110.
- Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name when configuring the service group.
See [“Setting the internal name of the clustered instance”](#) on page 116.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.

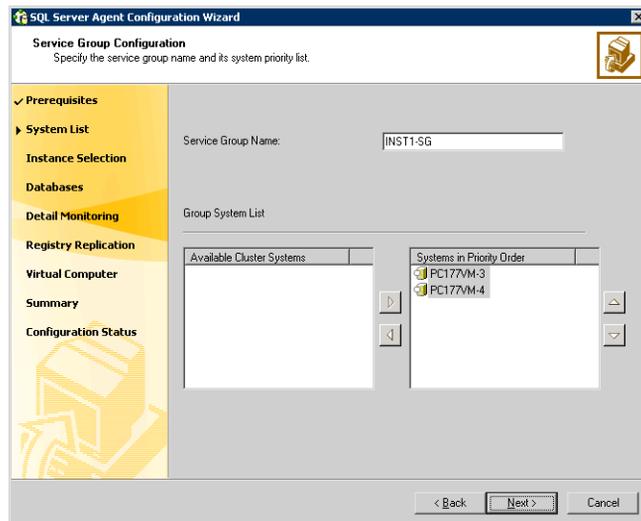
- Assign a unique virtual IP address for the SQL Server 2000 instance. You specify this IP address when configuring the service group.
- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists on all the cluster nodes. A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2000\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2000 Server service for the SQL instance. See [“Stopping the SQL Server 2000 service”](#) on page 109.

Creating the SQL Server 2000 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the

systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.

For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

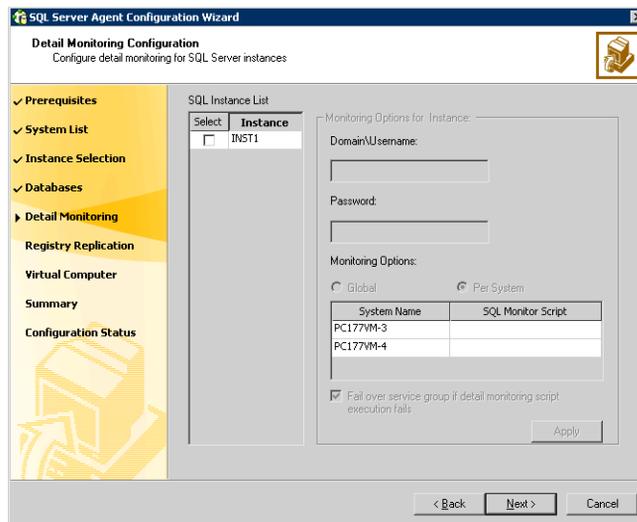
- Click **Next**.

- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.

Click **Next**.

- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

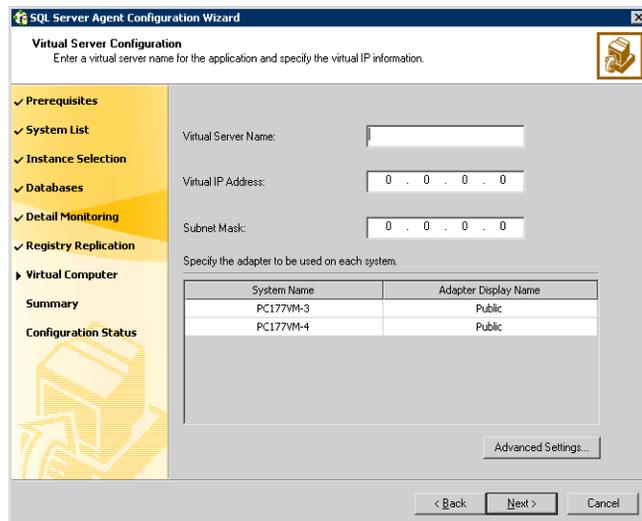
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 124 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script

on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.

- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 124 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.

- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
 - Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.
- The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

If you have created a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

See [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 133.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for high availability”](#) on page 355.

Verifying the SQL Server 2000 cluster configuration

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in step 1.
- 3 To move all the resources back to the original node, repeat step 1 for each of the service groups.

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in step 1.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.

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

Determining additional steps needed

Determine the next step as follows:

- If you have no more instances of SQL Server that you plan to install and configure for high availability, your high availability configuration is complete.
- If you have additional instances of SQL Server to install and configure in the cluster, start the installation and configuration sequence again for the next instance.
 - If you are using the Solutions Configuration Center:
See [“Following the workflow in the Solutions Configuration Center”](#) on page 69.
 - If you are using the printed document as a guide:
See [“Tasks for a new HA installation of SQL Server 2000”](#) on page 50.
- If you are planning to set up a secondary site for disaster recovery:
See [Chapter 16, “Deploying disaster recovery: New SQL Server 2000 installation”](#) on page 779.
- If you want to configure a Cluster Management Console connection:
See [“Configuring the Cluster Management Console connection”](#) on page 129
- If you need to modify the SQL Server service group:
See [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 133.
- If you need to make other changes to the cluster:
See *Veritas Cluster Server Administrator’s Guide*.

Configuring the Cluster Management Console connection

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

CMC comprises of the following components:

- *Management Server*

The management server accepts and processes the operational commands and the configuration inputs that users enter through CMC. The management server communicates with the VCS High Availability engine (HAD). Install the CMC Management Server only if you plan to centrally manage multiple clusters. You must install the management server on a standalone system that is outside any cluster but available on the local network.

- *Cluster Connector*

The cluster connector is an agent that enables the management server to communicate with clusters through intervening firewalls. You must install the cluster connector on each cluster that is separated from the management server by a firewall. If there are no firewalls between the management server and the clusters, you can configure the clusters to use direct connection instead.

In each cluster, the cluster connector runs on one node at a time, but is installed on all nodes and is configured for failover.

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

Prerequisites for installing the cluster connector

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

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

Installing the cluster connector on Windows clusters

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

To install the cluster connector on a Windows cluster

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

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

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

Files\Symantec\VRTScmccc\log

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

Configuring the cluster connector

Perform the following steps to configure the cluster connector.

To configure the cluster connector

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

Enter the domain type [1]: 4
- 9 Enter the password.
- 10 Enter a unique identifier for the cluster:

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

This indicates the configuration is successful.

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

Configuring the cluster connector using the management server console

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

To upgrade the cluster connector on discovered clusters

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

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

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

Modifying the SQL 2000 service group to add VMDg and MountV resources

If you create a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.

- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for high availability](#)” on page 355.

Deploying SFW HA for high availability: Standalone SQL Servers

This chapter includes the following topics:

- [“Tasks for converting a standalone SQL 2000 Server for high availability”](#) on page 136
- [“Reviewing the requirements”](#) on page 139
- [“Reviewing the configuration”](#) on page 143
- [“Configuring the storage hardware and network”](#) on page 145
- [“Preparing the standalone SQL Server”](#) on page 148
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 149
- [“Configuring cluster disk groups and volumes for SQL Server 2000”](#) on page 155
- [“Configuring the cluster”](#) on page 165
- [“Verifying that SQL Server 2000 databases and logs are moved to shared storage”](#) on page 181
- [“Preparing to install SQL Server 2000 on additional nodes”](#) on page 181
- [“”](#) on page 184
- [“Completing configuration steps in SQL Server”](#) on page 190
- [“Configuring the VCS SQL Server 2000 service group”](#) on page 191
- [“Verifying the SQL Server 2000 cluster configuration”](#) on page 197
- [“Adding a SQL Server user-defined database”](#) on page 198

- “[Additional instructions for disaster recovery](#)” on page 198

Tasks for converting a standalone SQL 2000 Server for high availability

You can convert a standalone SQL 2000 Server into a “clustered” SQL Server in a new Veritas Storage Foundation HA environment. This environment involves an Active-Passive configuration with one to one failover capabilities.

To plan a new SQL Server 2000 deployment, or to review additional considerations for an Active-Active configuration, refer to “[Deploying SFW HA for high availability: New SQL Server 2000 installation](#)” on page 49.

Note: In addition to the information contained in this chapter, the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach are required. Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Some installation and configuration options are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for wide area disaster recovery.

Symantec recommends using the Solutions Configuration Center as a guide for configuring high availability for SQL Server 2000.

See [Chapter 2, “Using the Solutions Configuration Center](#)” on page 31.

To configure MSDTC service groups, see [Chapter 8, “Configuring an MSDTC service group for high availability](#)” on page 355.

[Table 5-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 5-1 Tasks for converting a standalone SQL 2000 server for high availability

Objective	Tasks
“ Reviewing the requirements ” on page 139	■ Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 143	■ Understanding active-passive configuration ■ Reviewing the sample configuration

Table 5-1 Tasks for converting a standalone SQL 2000 server for high availability (Continued)

Objective	Tasks
“Configuring the storage hardware and network” on page 145	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Preparing the standalone SQL Server” on page 148	<ul style="list-style-type: none"> ■ Backing up existing data ■ Setting SQL Server services to manual start
“Installing Veritas Storage Foundation HA for Windows” on page 149	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2000” on page 155	<ul style="list-style-type: none"> ■ Planning the storage layout ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ For a new shared storage configuration, creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Configuring the cluster” on page 165	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“Verifying that SQL Server 2000 databases and logs are moved to shared storage” on page 181	<ul style="list-style-type: none"> ■ Stopping SQL Server service ■ Verifying that existing data is backed up ■ Modifying data file and user database locations ■ Restarting SQL Server service
“Preparing to install SQL Server 2000 on additional nodes” on page 181	<ul style="list-style-type: none"> ■ Ensuring that the cluster disk group is imported on the additional node ■ Optionally, renaming system data files

Table 5-1 Tasks for converting a standalone SQL 2000 server for high availability (Continued)

Objective	Tasks
" on page 184	<ul style="list-style-type: none"><li data-bbox="639 348 1086 374">■ Installing SQL Server on additional nodes<li data-bbox="639 380 1175 406">■ Setting the internal name of the clustered instance
"Configuring the VCS SQL Server 2000 service group" on page 191	Creating a SQL Server service group using the VCS SQL Configuration wizard
"Verifying the SQL Server 2000 cluster configuration" on page 197	<ul style="list-style-type: none"><li data-bbox="639 588 872 614">■ Simulating failover<li data-bbox="639 619 911 645">■ Switching online nodes

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at:

<http://www.symantec.com/business/support/index.jsp>

Supported Software

- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the Global Clustering Option and depending on your replication solution, Veritas Volume Replicator or a hardware replication agent

For a Microsoft SQL Server 2000 environment, any of the following SQL Servers and their operating systems:

- | | |
|---|---|
| <p>Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) |
| <p>Microsoft SQL Server 2000 (64-bit) Enterprise Edition</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| <p>Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions) |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 142.
- 1 GB of RAM per server for SFW HA.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL Virtual Server

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

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on `C:\WINDOWS` of one node, installations on all other nodes must be on `C:\WINDOWS`. Make sure that the

same drive letter is available on all nodes and that the system drive has adequate space for the installation.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Disk space requirements

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

[Table 5-2](#) on page 142 estimates disk space requirements for SFW HA.

Table 5-2 Disk space requirements

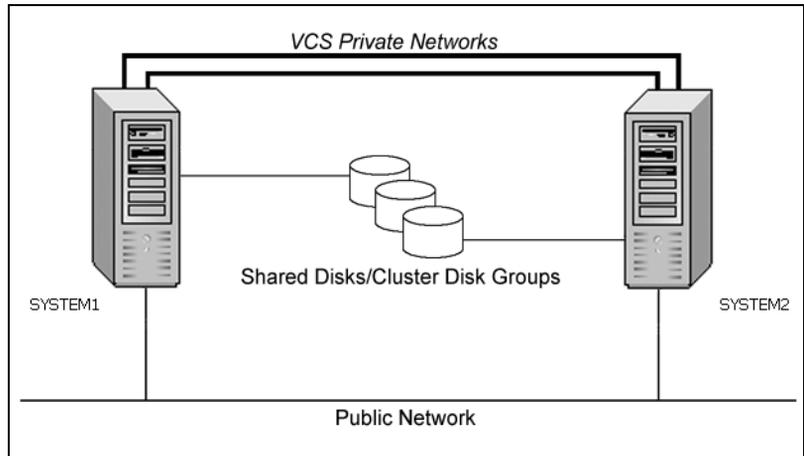
Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Reviewing the configuration

In a typical example of a high availability cluster, you create a virtual server in an Active-Passive SQL configuration. 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.

Figure 5-1 illustrates a typical Active-Passive configuration. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured on the active node (SYSTEM1). If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 5-1 Active-Passive configuration



The virtual SQL Server is online on SYSTEM1, serving client requests. The shared disks provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. From the user's perspective there will be a small delay as the backup node comes online, but the interruption in effective service is minimized.

Sample configuration

A sample setup is used to illustrate the installation and configuration tasks. During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the IP address should be the same on all nodes
- Cluster IP address: used by Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console

You should have these IP addresses available before you start deploying your environment. The following names describe the objects created and used during the installation and configuration:

Table 5-3 Standalone SQL Server configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL Server
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1	SQL Instance Name
INST1-VS	name of the SQL Virtual Server

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Preparing the standalone SQL Server

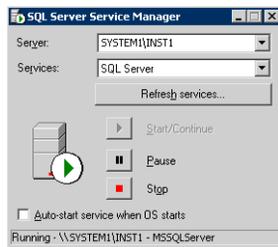
Before you begin the process of installing Veritas Storage Foundation HA for Windows and creating a clustered environment, be sure to back up SQL data on the existing standalone SQL Server.

In addition, complete the following task:

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 5-4](#) on page 149 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 5-4 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 150.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

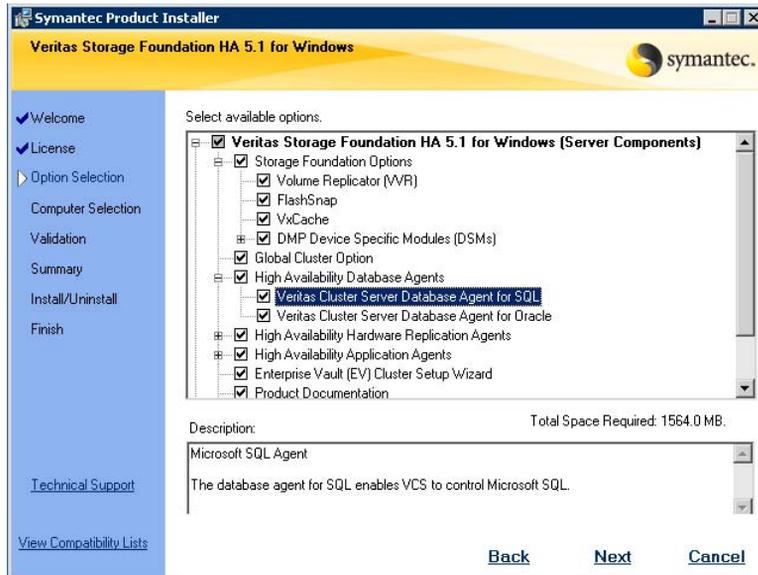
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

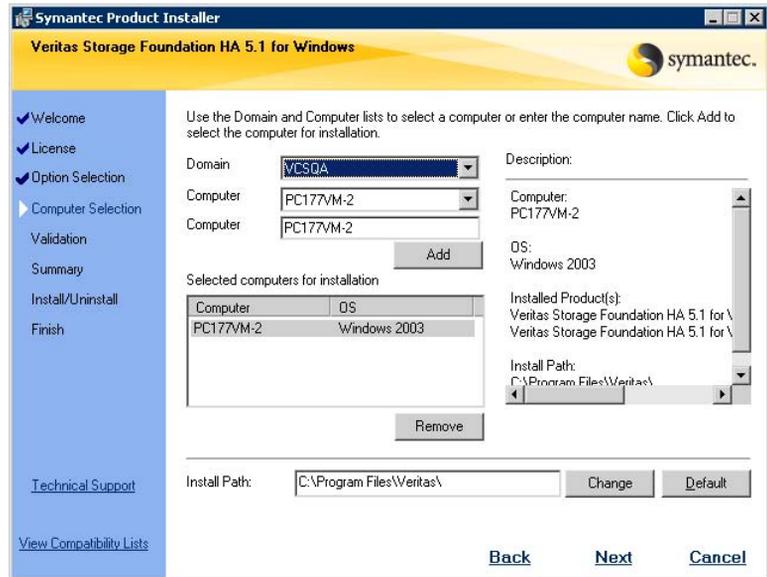
Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option Required for a disaster recovery configuration only.

Veritas Volume Replicator For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

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

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Configuring cluster disk groups and volumes for SQL Server 2000

You create cluster disk groups and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [“About cluster disk groups and volumes”](#) on page 156
- [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 156
- [“Sample disk group and volume configuration”](#) on page 157
- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 158
- [“Considerations for volumes for a VVR disaster recovery configuration”](#) on page 158
- [“Creating a cluster disk group”](#) on page 159
- [“Creating volumes”](#) on page 161

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: Create a cluster disk group and volumes for the SQL instance on only one node of a cluster. You make the volumes accessible 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes for the SQL instance:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
See “[Sample disk group and volume configuration](#)” on page 157.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.

- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

You may be configuring new shared storage for the high availability environment, or the existing standalone SQL Server databases and logs may already be on shared storage. If the existing databases and logs are already on shared storage, read the following topic:

- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 158

For a disaster recovery configuration using Veritas Volume Replicator, read the following topic:

- [“Considerations for volumes for a VVR disaster recovery configuration”](#) on page 158

Sample disk group and volume configuration

You first create a cluster disk group (INST1_DG) on shared disks. You then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process:

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

Your configuration may require additional planning. See the following topics:

- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 158
- [“Considerations for volumes for a VVR disaster recovery configuration”](#) on page 158

Considerations for converting existing shared storage to cluster disk groups and volumes

The databases and logs for your existing standalone SQL Server may already be on shared storage. In this case, when you create cluster disk groups, you specify the disks that contain the existing databases and logs.

Creating a disk group converts the disks from basic disks to dynamic disks. Partitions on the disks are automatically converted to volumes on the dynamic disks.

Therefore, if your existing disk layout contains databases and logs in the same partition, they become part of the same volume in the cluster disk group. If the disk contains multiple partitions, each containing a user database, each partition becomes a separate volume, but all will become part of the same cluster disk group. If this configuration does not meet your requirements, you may want to modify your disk layout before creating the cluster disk group.

For additional information on converting basic to dynamic disks, see *Veritas Storage Foundation Administrator's Guide*.

Symantec recommends creating a separate 100 MB RegRep volume that contains the list of registry keys that must be replicated among cluster systems for the SQL service. However, if no additional disks are available on the shared storage, you can specify an existing volume as the registry replication path during service group creation.

For a disaster recovery configuration using Veritas Volume Replicator, you will need to allow additional disk space for a Storage Replicator Log volume.

See "[Considerations for volumes for a VVR disaster recovery configuration](#)" on page 158.

Considerations for volumes for a VVR disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- When you run the Disaster Recovery Wizard, a VVR Storage Replicator Log (SRL) volume is created automatically for each disk group that contains volumes that are replicated. Ensure that you allow sufficient disk space for this volume. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.
- Symantec recommends that for replication considerations, you create a separate volume for tempdb, for example, INST1_TEMPDB, within the system database disk group. When you later configure replication for disaster recovery, you replicate that disk group but exclude the tempdb volume from the replication.

It would waste bandwidth to replicate tempdb because the data is transitory and is not needed for DR site recovery.

You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume. See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 190.

- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Creating a cluster disk group

Use the Veritas Enterprise Administrator (VEA) to create a cluster disk group on the existing standalone SQL Server system. Repeat the procedure if you want to create additional disk groups.

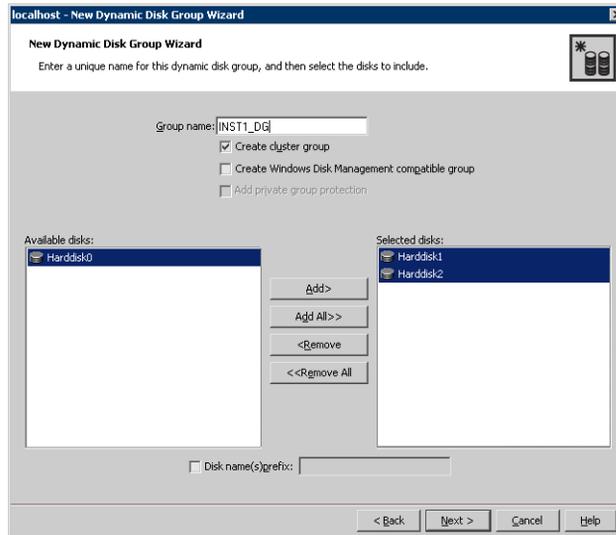
Ensure that you have reviewed all prerequisites and considerations before creating the cluster disk group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This procedure will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure to create additional volumes. For example, create the following volumes on the INST1-DG disk group:

- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL**: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

You could create the following volumes on the same disk group as the system volumes or on a separate disk group, depending on how you prefer to configure your storage:

- **INST1_DB1_VOL**: contains the user database files
- **INST1_DB1_LOG**: contains the user database log files

Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.

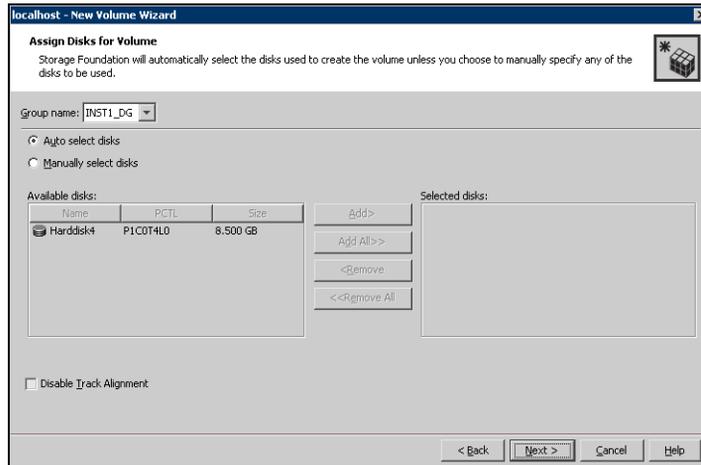
See “[Considerations for volumes for a VVR disaster recovery configuration](#)” on page 158.

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

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

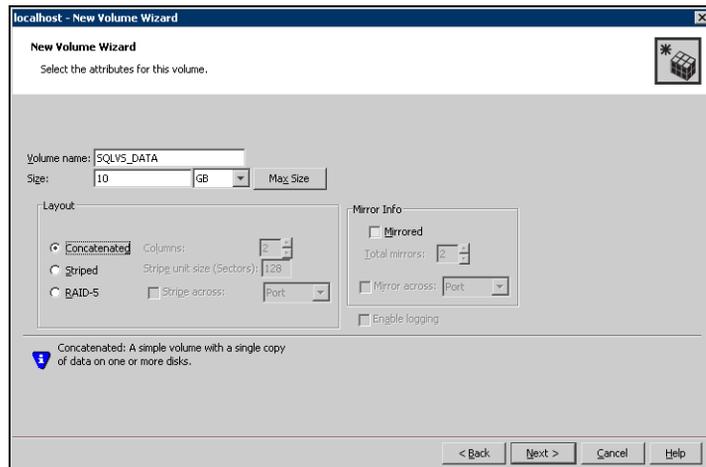


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

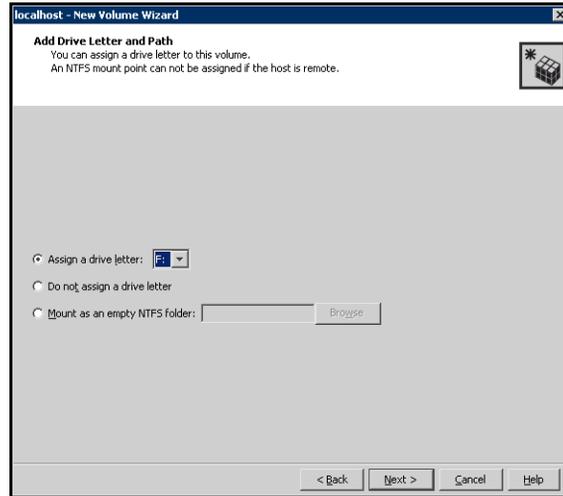
- 8 Click **Next**.

9 Specify the parameters of the volume.



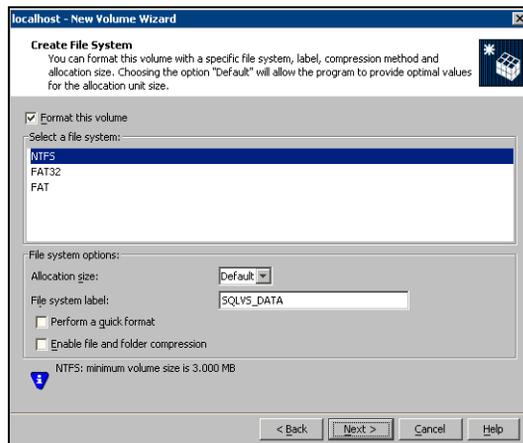
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the **Default**.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Configuring the cluster

You configure the cluster to contain the existing SQL Server 2000 system and the system that will become the failover node. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also provides the option to configure the ClusterService group, which can contain resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

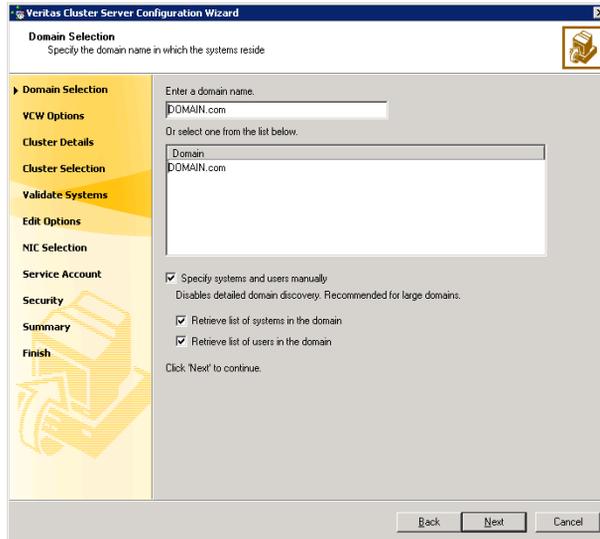
- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 139.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or removing or modifying cluster configurations.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



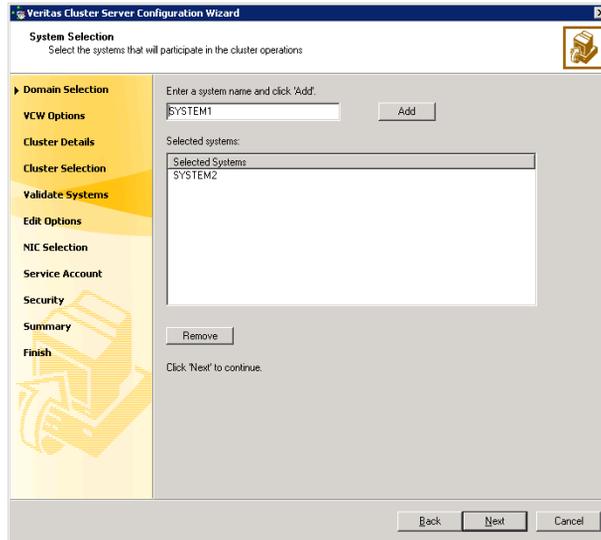
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 8](#) on page 168.

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

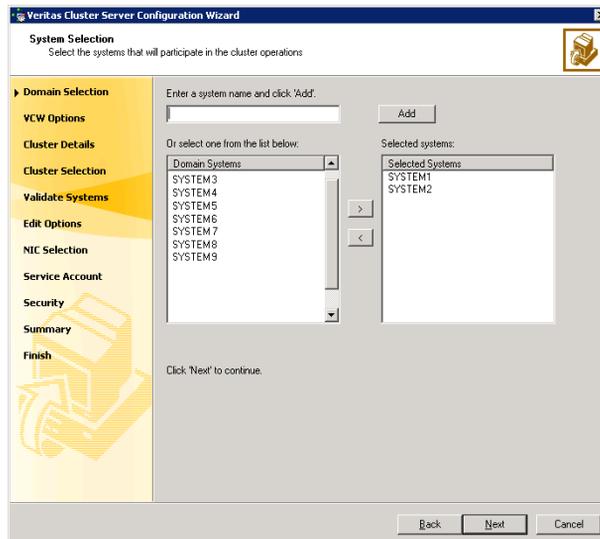
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 167. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 168.

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



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

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

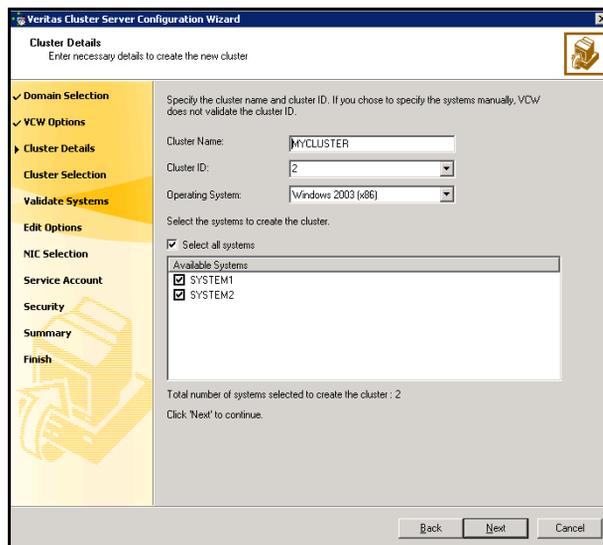
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 166 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

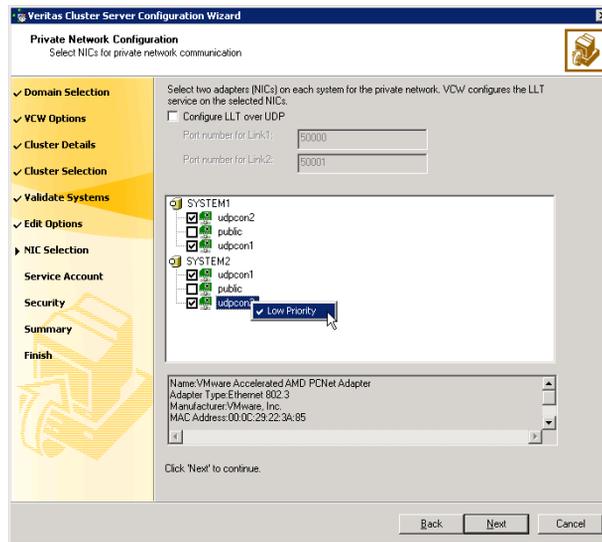
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 168, proceed to the next step. Otherwise, proceed to [step 12](#) on page 172.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

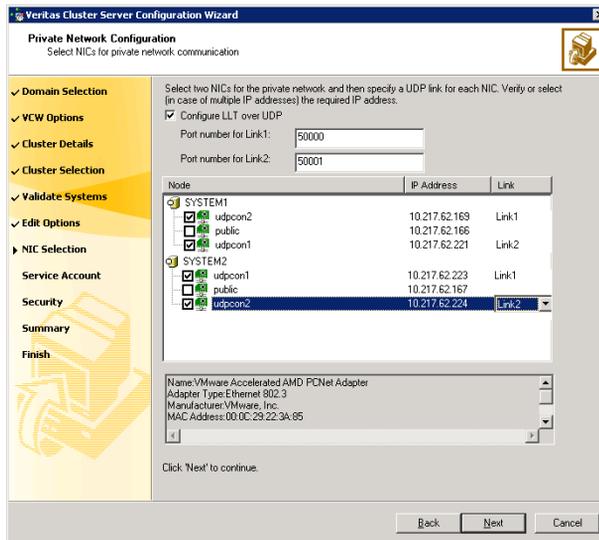
- To configure the VCS private network over Ethernet



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

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

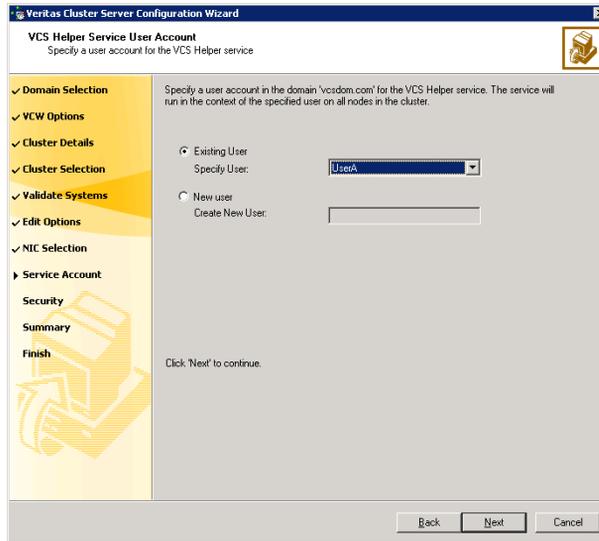
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 166, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

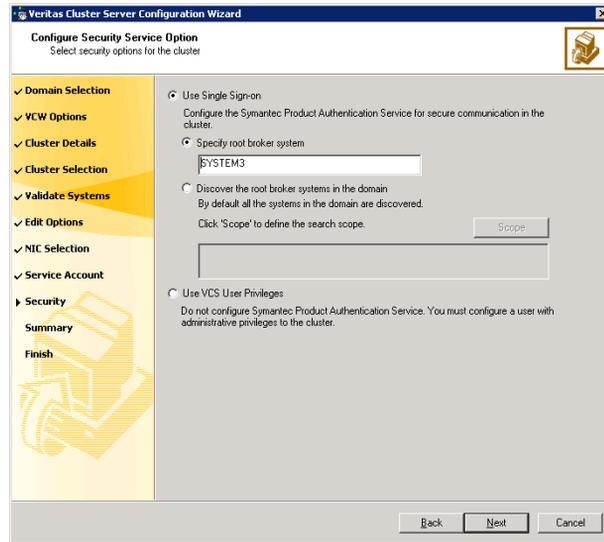
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

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

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 5-5](#) on page 174 contains some more examples of search criteria.

Table 5-5 Search criteria examples

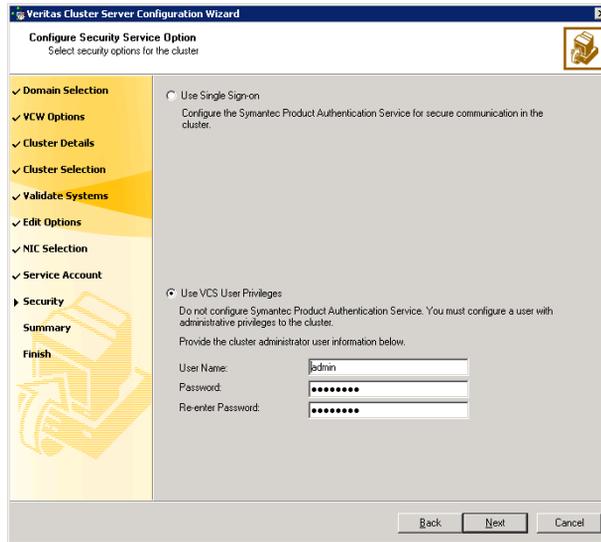
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

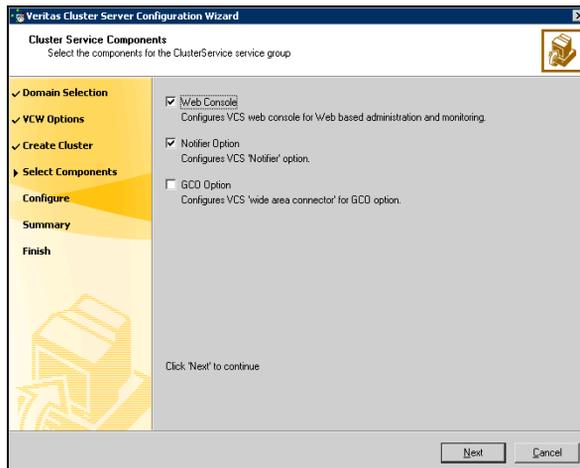
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 177.

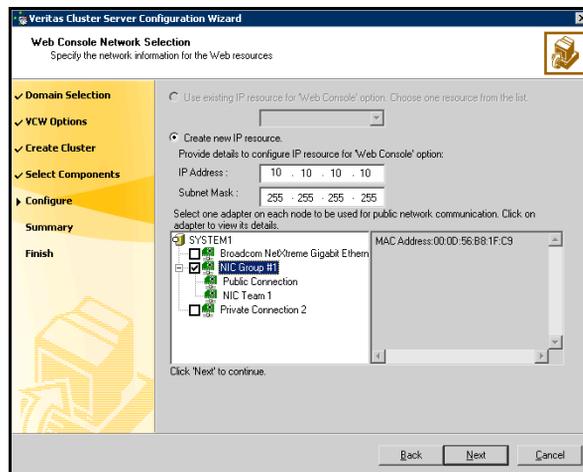
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 178.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



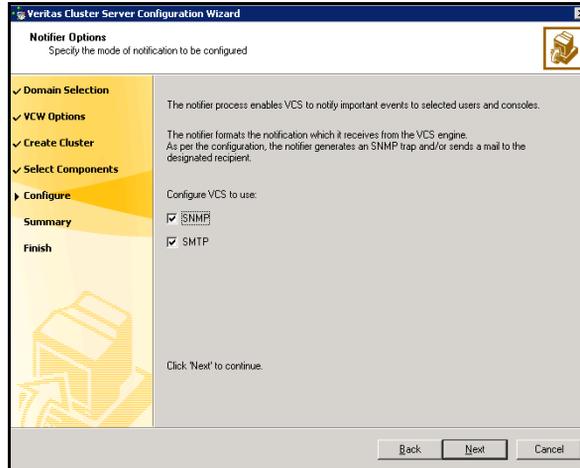
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 178. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

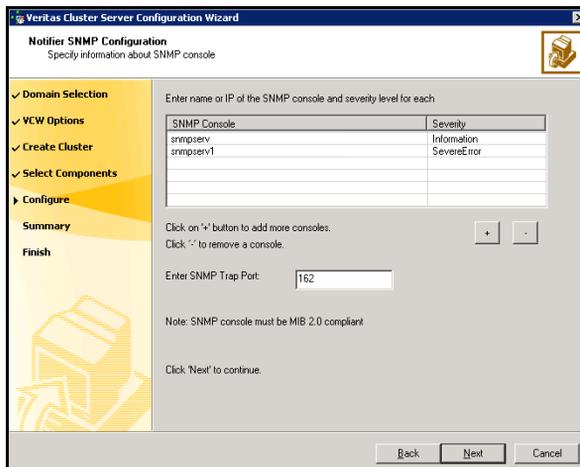
To configure notification

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

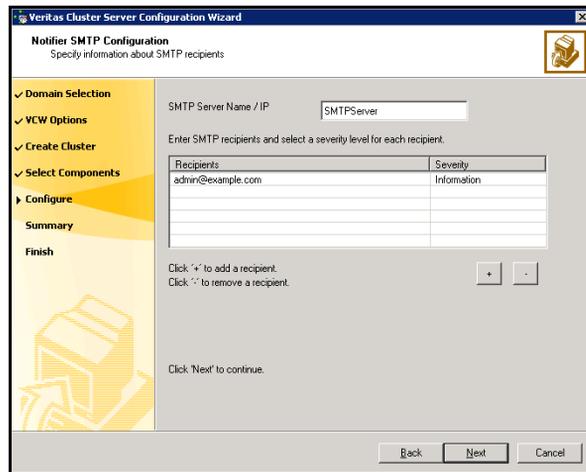


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

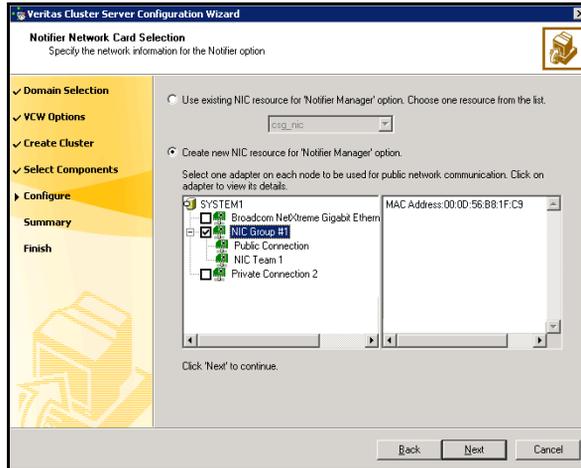


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Verifying that SQL Server 2000 databases and logs are moved to shared storage

Verify the location of all SQL Server databases and logs for the existing standalone server. If they are located on local storage, move them from the local drive to the appropriate cluster disk groups and volumes on shared storage to ensure proper failover operations in the cluster.

Complete the following tasks to move the databases:

- 1 Stop the SQL Server service.
- 2 Verify that you have backed up your existing data.
- 3 Ensure that the dynamic disk group is imported on the node where the original files are located on the local drives and mount the volumes. Refer to “[Deporting the cluster disk group](#)” on page 182, “[Importing the cluster disk group](#)” on page 182, and “[Adding drive letters to mount the volumes](#)” on page 183 for instructions.
- 4 Modify the SQL Server 2000 Data File and User Database locations. Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.
<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>
- 5 Restart SQL Server 2000.

Preparing to install SQL Server 2000 on additional nodes

Ensure that the cluster disk group is imported on the node where you install SQL Server 2000. In addition, you may want to rename the system data files before installation, so that they can be restored if anything goes wrong with the installation.

See the following topics:

- “[Deporting the cluster disk group](#)” on page 182
- “[Importing the cluster disk group](#)” on page 182
- “[Adding drive letters to mount the volumes](#)” on page 183
- “[Renaming shared SQL Server 2000 files](#)” on page 184

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

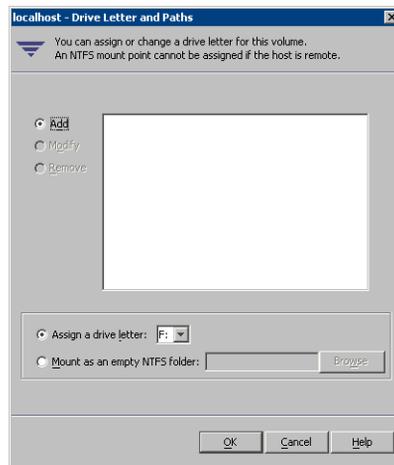
- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
 Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2000 on additional nodes

Install Microsoft SQL Server on each additional node that will be included in the SQL Server service group. Make sure to install the SQL Server data files to shared storage.

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2000 on additional nodes”](#) on page 184
- [“Setting the internal name of the clustered instance”](#) on page 188

Installing SQL Server 2000 on additional nodes

Before installing Microsoft SQL Server 2000, complete the preparation steps.

See [“Preparing to install SQL Server 2000 on additional nodes”](#) on page 181.

SQL Server 2000 must have the same configuration on all nodes in the cluster. You will need the following information:

- Instance name (if applicable)
- Destination folder for program files and data files
- Authentication mode

Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Multiple instances of SQL Server must be installed in the same order on every node of the cluster.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

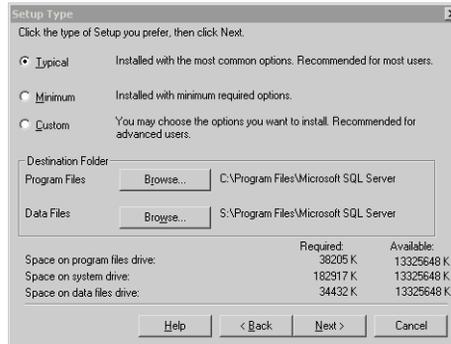
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

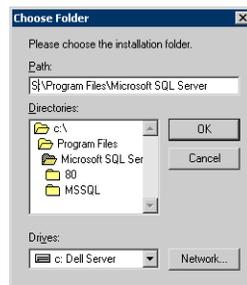
- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.

Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.

- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

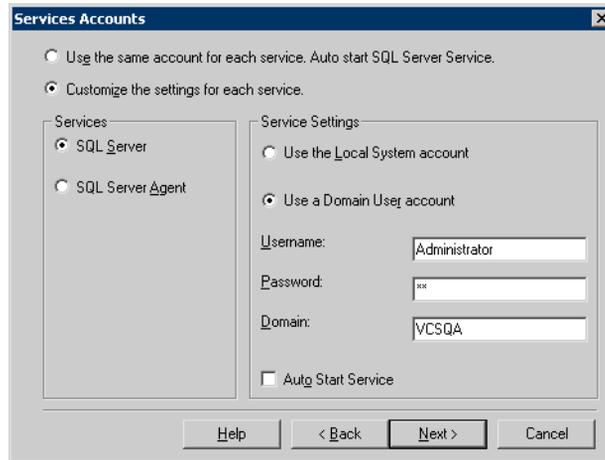


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
- In the Services box, select the **SQL Server** option.
- In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
- Clear the **Auto Start Service** option.
- Repeat these steps for the SQL Server Agent option.

11 Follow the wizard instructions to complete the installation.

12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in [“Preparing to install SQL Server 2000 on additional nodes”](#) on page 181 and [“”](#) on page 184 on any additional nodes.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

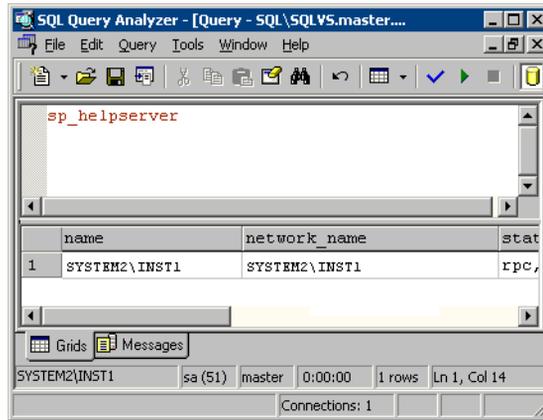
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example *SYSTEM2\INST1*.
- Select the **Start SQL server if it is stopped checkbox**.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
 “sp_dropserver `System_Name\Instance_Name.`”
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
 “sp_dropserver `SYSTEM2\INST1.`”
 For example, for a default instance:
 “sp_dropserver `SYSTEM1.`”
 - Press F5.
- 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter
`"sp_addserver `Virtual_Server_Name\Instance_Name`, local"`
For example ``INST1-VS\INST1``, `local` for a named instance, or ``INST1-VS``, `local` for a default instance.
 - Press F5.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume.

See "[Moving the tempdb database if using VVR for disaster recovery](#)" on page 190.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See "[Assigning the port for multiple instances](#)" on page 191.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See "[Creating volumes](#)" on page 161.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Assigning the port for multiple instances

When you are running multiple SQL Server instances, you must assign a different port to each instance.

Refer to the Microsoft Knowledge Base for the instructions on assigning static ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

Configuring the VCS SQL Server 2000 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

A VCS SQL Server service group is used to bring a SQL Server 2000 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Read the following topics:

- [Prerequisites for configuring the service group](#)
- [Creating the SQL Server 2000 service group](#)

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See “[Installing Veritas Storage Foundation HA for Windows](#)” on page 149.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See “[Configuring the cluster](#)” on page 165.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.

- Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group. See “[Importing the cluster disk group](#)” on page 182.
- Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- To avoid having to modify the service group later, create any new user-defined databases before creating the service group. See “[Adding a SQL Server user-defined database](#)” on page 198.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name when configuring the service group. See “[Setting the internal name of the clustered instance](#)” on page 188.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.

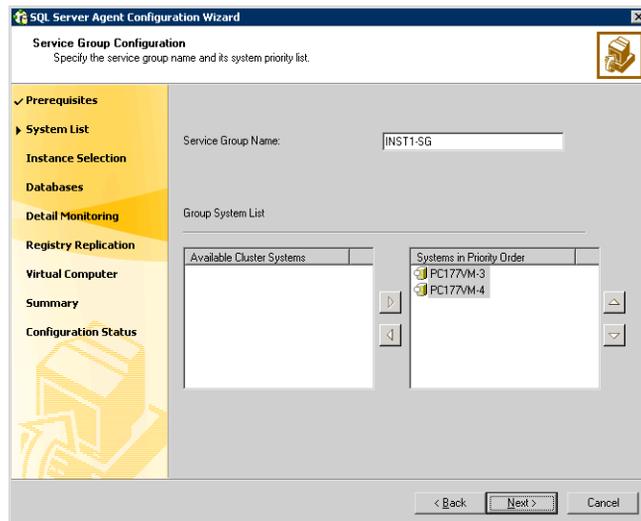
- Assign a unique virtual IP address for the SQL Server 2000 instance. You specify this IP address when configuring the service group.
- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists on all the cluster nodes. A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2000\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2000 Server service for the SQL instance.

Creating the SQL Server 2000 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the

systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.

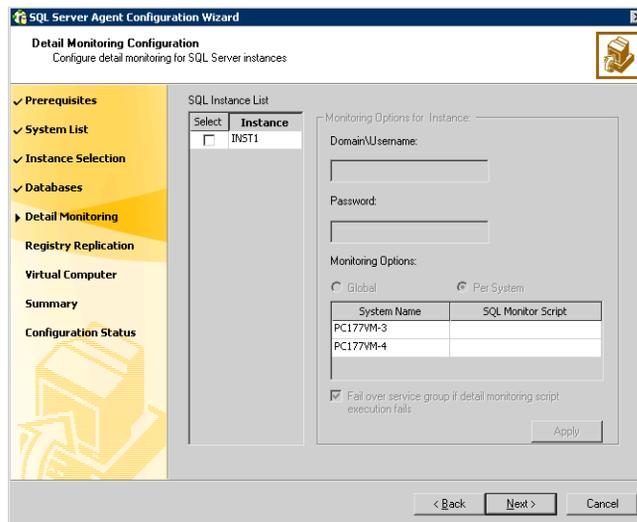
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

■ Click **Next**.

- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.

Click **Next**.

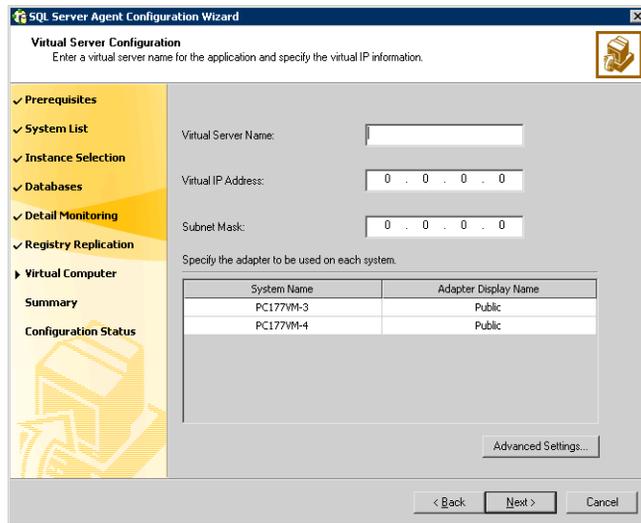
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 194 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the

script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.

- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 194 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example **INST1-VS**. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.

- Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for high availability](#)” on page 355.

Verifying the SQL Server 2000 cluster configuration

Verify your installation by switching online nodes or by shutting down the computer that is currently online. Either process will test that the service group can be smoothly transferred between nodes.

Shutting down a node creates an actual failure, stressing your system, but more truly testing its high availability than by switching nodes. If you do shut down the online computer in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another node.

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

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

To switch service groups

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

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.
- 4 If you need to move all the service groups back to the original node:

- Restart the node you shut down in step 1.
- Click **Switch To**, and click the appropriate node from the menu.
- In the dialog box, click **Yes**.
The service group you selected is taken offline and brought online on the node that you selected.

Adding a SQL Server user-defined database

If you want to add a SQL Server user-defined database to an existing configuration, complete the procedure “[Creating a SQL Server user-defined database](#)” on page 118.

The procedure includes the following tasks:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Additional instructions for disaster recovery

After completing the tasks for setting up a SFW HA environment for SQL Server 2000 on a primary site, you can optionally create a secondary or “failover” site for disaster recovery.

See [Chapter 16, “Deploying disaster recovery: New SQL Server 2000 installation”](#) on page 779.

Deploying SFW HA for high availability: New SQL Server 2005 installation

This chapter covers the following topics:

- [“Tasks for a new HA installation of SQL Server 2005”](#) on page 200
- [“Reviewing the requirements”](#) on page 205
- [“Reviewing the configuration”](#) on page 209
- [“Reviewing considerations for Active-Active configurations”](#) on page 217
- [“Configuring the storage hardware and network”](#) on page 220
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 223
- [“Configuring cluster disk groups and volumes for SQL Server 2005”](#) on page 230
- [“Configuring the cluster”](#) on page 239
- [“About installing multiple instances”](#) on page 256
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 256
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 261
- [“Installing and configuring SQL Server 2005 on the second node”](#) on page 264
- [“Creating a SQL Server user-defined database”](#) on page 271
- [“Completing configuration steps in SQL Server”](#) on page 272

- [“Configuring the VCS SQL Server 2005 service group”](#) on page 274
- [“Verifying the SQL Server 2005 cluster configuration”](#) on page 280
- [“Determining additional steps needed”](#) on page 281

Tasks for a new HA installation of SQL Server 2005

You can install and configure a new Veritas Storage Foundation HA environment for SQL Server 2005 in the following ways:

Active-Passive	<p>One SQL instance per node with one to one failover capabilities.</p> <p>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.</p>
Active-Active	<p>Multiple SQL instances per cluster node.</p> <p>For example, in a two-node cluster with two SQL 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.</p>

Note: Some installation and configuration options are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2005.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for high availability”](#) on page 355.

Tasks for an Active-Passive configuration

Table 6-1 outlines the high-level objectives and the tasks to complete each objective for an Active-Passive configuration.

Table 6-1 SQL Server 2005: Active-Passive configuration tasks

Objective	Tasks
“Reviewing the requirements” on page 205	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 209	<ul style="list-style-type: none"> ■ Understanding active/passive configuration ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 220	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Installing Veritas Storage Foundation HA for Windows” on page 223	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“Configuring cluster disk groups and volumes for SQL Server 2005” on page 230	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Configuring the cluster” on page 239	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“Installing and configuring SQL Server 2005 on the first node” on page 256	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2005 ■ Setting SQL Server services to manual start

Table 6-1 SQL Server 2005: Active-Passive configuration tasks (Continued)

Objective	Tasks
“Preparing to install SQL Server 2005 on the second node” on page 261	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2005 on the second node” on page 264	Installing and configuring SQL Server 2005 on the second and any additional nodes
“Setting the internal name of the clustered instance” on page 268	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 271	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 272	Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2005 service group” on page 274	Creating a SQL Server service group using the SQL Server Configuration Wizard
“Verifying the SQL Server 2005 cluster configuration” on page 280	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes

Tasks for an Active-Active configuration

Table 6-2 outlines the high-level objectives and the tasks to complete each objective for an Active-Active configuration.

Table 6-2 SQL Server 2005: Active-Active configuration tasks

Objective	Tasks
“ Reviewing the requirements ” on page 205	Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 209	<ul style="list-style-type: none"> ■ Understanding active-active configuration ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 220	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Installing Veritas Storage Foundation HA for Windows ” on page 223	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“ Configuring cluster disk groups and volumes for SQL Server 2005 ” on page 230	<ul style="list-style-type: none"> ■ Creating a dynamic disk group for each instance ■ Creating dynamic volumes for the SQL system database, user databases, transaction logs, and replicated registry keys for each instance
“ Configuring the cluster ” on page 239	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“ Installing and configuring SQL Server 2005 on the first node ” on page 256	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server on the first node ■ Setting SQL Server services to manual start

Table 6-2 SQL Server 2005: Active-Active configuration tasks (Continued)

Objective	Tasks
“Preparing to install SQL Server 2005 on the second node” on page 261	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Renaming or removing shared SQL files
“Installing and configuring SQL Server 2005 on the second node” on page 264	Installing SQL Server on the second node
“Setting the internal name of the clustered instance” on page 268	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 271	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Completing configuration steps in SQL Server” on page 272	Doing additional configuration steps in SQL Server for multiple instances or disaster recovery
“Configuring the VCS SQL Server 2005 service group” on page 274	<ul style="list-style-type: none"> ■ Creating a SQL Server service group using the VCS SQL Configuration Wizard ■ Ensuring the priority order of the systems is set up in reverse order for each instance
“Verifying the SQL Server 2005 cluster configuration” on page 280	<ul style="list-style-type: none"> ■ Simulating fail over ■ Switching online nodes
“Determining additional steps needed” on page 281	Repeating the installation and configuration steps for the next instance

Reviewing the requirements

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation.

Disk space requirements

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

[Table 6-3](#) on page 205 estimates disk space requirements for SFW HA.

Table 6-3 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 for 64-bit Itanium (IA64)■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 209.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW

HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

Review the information for the configurations you have planned:

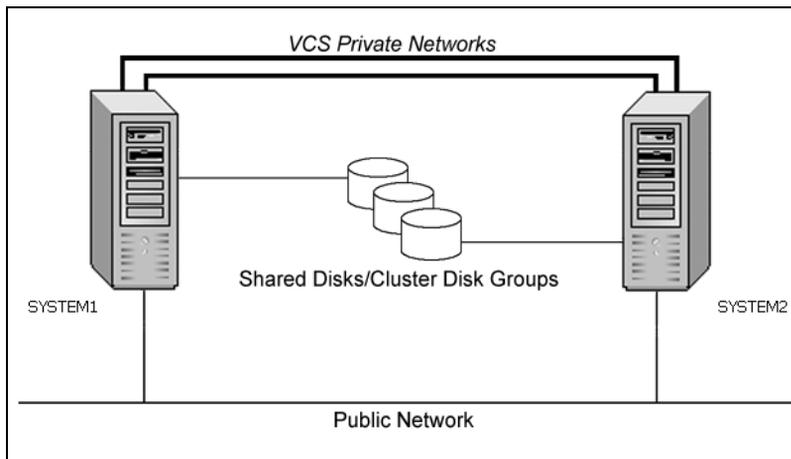
- [Active-Passive configuration](#)
- [Active-Active configuration](#)
- [Disaster recovery configuration](#)

Active-Passive configuration

In a typical example of a high availability cluster, you create a virtual SQL server in an Active-Passive configuration. 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.

Figure 6-1 illustrates a typical Active-Passive configuration. SQL Server is installed on both SYSTEM1 and SYSTEM2 and configured as a virtual server (INST1-VS) with a virtual IP address. The SQL databases are configured on the shared storage on volumes contained in cluster disk groups. The SQL virtual server is configured to come online on SYSTEM1 first. If SYSTEM1 fails, SYSTEM2 becomes the active node and the SQL virtual server comes online on SYSTEM2.

Figure 6-1 Active-Passive configuration



Sample Active-Passive configuration

A sample setup is used to illustrate the installation and configuration tasks.

The following names describe the objects created and used during the installation and configuration:

Table 6-4 Active-Passive configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
INST1_SG	SQL Server service group

IP addresses for sample Active-Passive configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there is one SQL Server virtual server. Therefore you would need one virtual server IP address. If you want to use the VCS Web Console or the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

See “Network requirements” under “[Reviewing the requirements](#)” on page 205.

See “[Disaster recovery configuration](#)” on page 214.

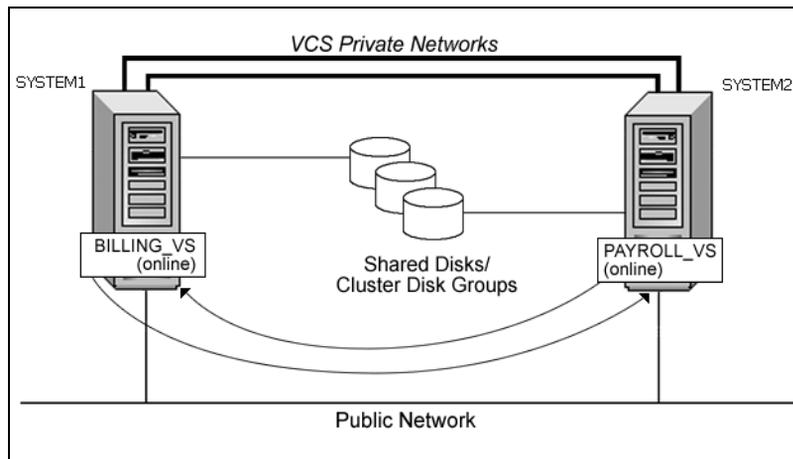
Active-Active configuration

In an Active-Active SQL Server configuration, several instances are intended to run on a single node when necessary. A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

A SQL Server instance can fail over to any of the other cluster nodes that you specify when you configure the SQL Server service group.

The following figure illustrates a two node Active-Active configuration. The SQL Server databases are configured on the shared storage on volumes contained in cluster disk groups. Each SQL Server virtual server is configured in a separate SQL Server service group. Each service group can fail over to the other node in the cluster.

Figure 6-2 Active-Active configuration



For example, consider a two-node cluster hosting two SQL Server virtual servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

SQL Virtual Server	Service Group	System List
BILLING_VS	BILLING_SG	SYSTEM1, SYSTEM2
PAYROLL_VS	PAYROLL_SG	SYSTEM2, SYSTEM1

Sample Active-Active configuration

A sample setup is used to illustrate the installation and configuration tasks for two instances of SQL Server, Billing and Payroll. During normal operation, one instance will be online on each of the two servers. If a failure occurs, the instance on the failing node will be brought online on the other server, resulting in two instances running on one server.

The following names describe the objects created and used during the installation and configuration:

Table 6-5 Active-Active configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	server names
BILLING_DG	cluster disk group for the billing instance
PAYROLL_DG	cluster disk group for the payroll instance
BILLING_VS_SYS_FILES	volume for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume for the SQL Server system data files for the payroll instance
BILLING_DATA	volume for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume for a SQL Server user-defined database log file for the billing instance
PAYROLL_LOG	volume for a SQL Server user-defined database log file for the payroll instance
BILLING_REGREP	volume for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume for the list of registry keys replicated among the nodes for the payroll instance

Table 6-5 Active-Active configuration objects

Object Name	Description
SQL_CLUS1	virtual SQL Server cluster
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL Server name for the billing instance
PAYROLL_VS	virtual SQL Server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

IP addresses for sample Active-Active configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration. Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there are two virtual servers: BILLING-VS and PAYROLL-VS. Therefore, you would need two virtual server IP addresses. If you want to use the VCS Web Console or the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

See “Network requirements” under “[Reviewing the requirements](#)” on page 205.

See “[Disaster recovery configuration](#)” on page 214.

Disaster recovery configuration

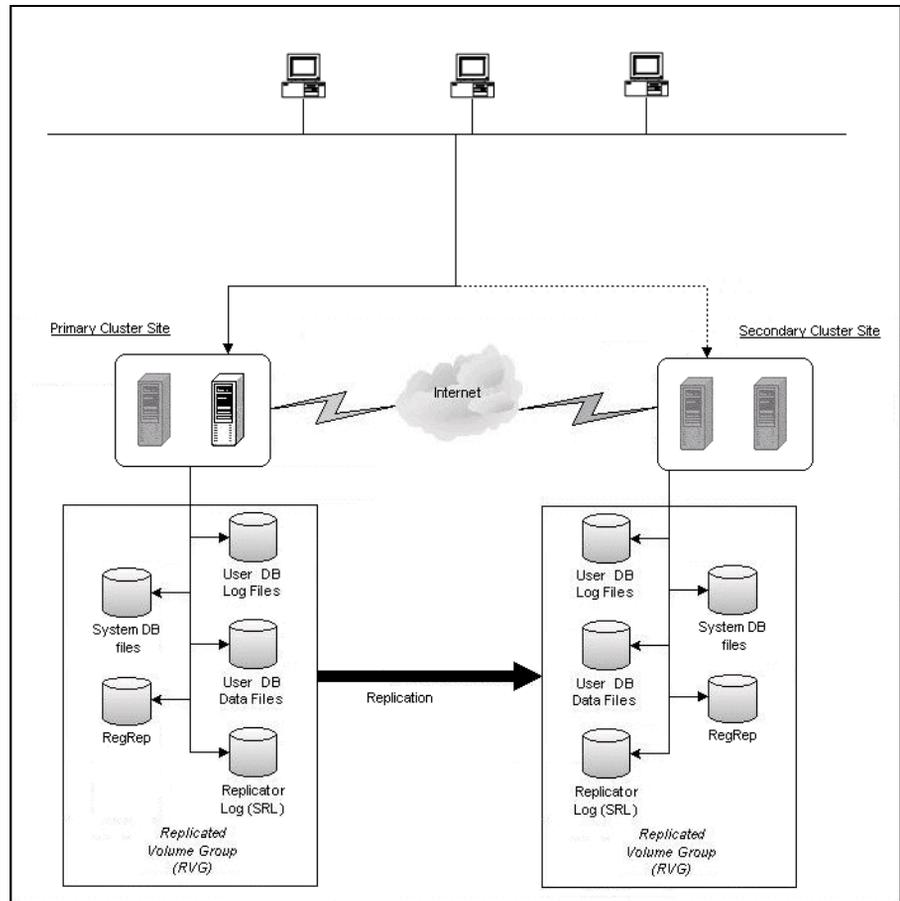
You may be preparing to configure both a primary site and a secondary site for disaster recovery.

Figure 6-3 on page 215 illustrates a typical Active-Passive disaster recovery configuration using Veritas Volume Replicator (VVR).

In the example illustration, the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the Replicated Volume Group (RVG). The SQL Server application data is stored on the volumes that are under the control of the RVG.

If the Microsoft SQL Server server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Figure 6-3 Typical Active-Passive disaster recovery configuration with VVR



Sample disaster recovery configuration

The sample disaster recovery setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

A sample setup is used to illustrate the installation and configuration tasks:

Table 6-6 Disaster Recovery configuration objects

Object Name	Description
Primary Site	
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
Secondary Site	
SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.
DR Components	
INST1_REPLOG	replicator log volume required by VVR
INST1_DB1_RDS	RDS (Replicated Data Set) name
INST1_DB1_RVG	RVG (Replicated Volume Group) name
INST1_DB1_RVG_SG	Replication service group

IP addresses for sample disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, you should obtain all required IP addresses before beginning configuration.

See “Network requirements” under “[Reviewing the requirements](#)” on page 205.

You specify the following addresses during the replication process:

Table 6-7 IP addresses required for DR configuration

IP Address	Description
SQL virtual server IP address	For a disaster recovery configuration, the virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.
Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.
Replication IP address	You need two IP addresses per application instance, one for the primary site and one for the secondary site.

Reviewing considerations for Active-Active configurations

For an Active-Active configuration or in other cases when you are installing multiple instances of SQL Server on the same system, there are special considerations.

To assist you in planning your deployment, the considerations are summarized in the following topics:

- “[Key information for Active-Active configurations](#)” on page 218
- “[Following the workflow in the Solutions Configuration Center](#)” on page 219

Key information for Active-Active configurations

The following table summarizes key information about Active-Active configurations and multiple instances and cross-references additional information:

Table 6-8 Key information for Active-Active configuration

Task	Description
Configuring disk groups and volumes	<p>Create a separate set of cluster disk groups and volumes for each instance. You can create all the disk groups and volumes at one time or create them as a separate step for each instance.</p> <p>See “Considerations for disk groups and volumes for multiple instances” on page 232.</p>
Configuring the cluster	<p>If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. That way, you do not need to run the wizard again later to add the nodes.</p> <p>See “Configuring the cluster” on page 239.</p>
Installing and configuring SQL Server	<p>Assign a unique instance name, virtual server name, and port to each instance.</p> <p>“About installing multiple instances” on page 256.</p> <p>“Setting the internal name of the clustered instance” on page 268.</p> <p>“Assigning ports for multiple SQL Server instances” on page 273.</p>
Configuring the service group	<p>For an active/active configuration, create a separate service group for each instance. Each service group must have a unique name and virtual IP address. There are also special considerations for specifying the priority order of systems for failover.</p> <p>See “Service group requirements for active-active configurations” on page 274.</p>

Following the workflow in the Solutions Configuration Center

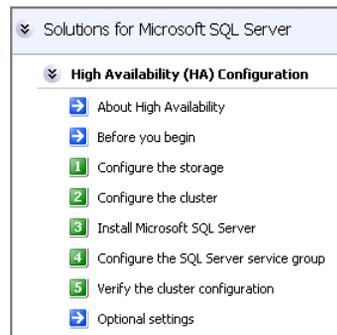
The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA environment for one or more instances of SQL Server 2005, in either an active-passive or active-active configuration.

Figure 6-4 on page 219 shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

If you are setting up multiple instances of SQL in the cluster, you may find it helpful to use the Configuration Center as follows:

- Under High Availability (HA) Configuration, complete all the steps for the first instance.
- For the next instance:
 - For step 1, Configure the storage: If you configured disk groups and volumes for the instance earlier, verify that they are available and continue with step 2.
 - For step 2, Configure the cluster: If you configured the nodes as part of the cluster earlier, as recommended, continue with step 3 and complete all subsequent steps.

Figure 6-4 Configuration steps in the Solutions Configuration Center



See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 6-9](#) on page 223 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 6-9 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 224.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

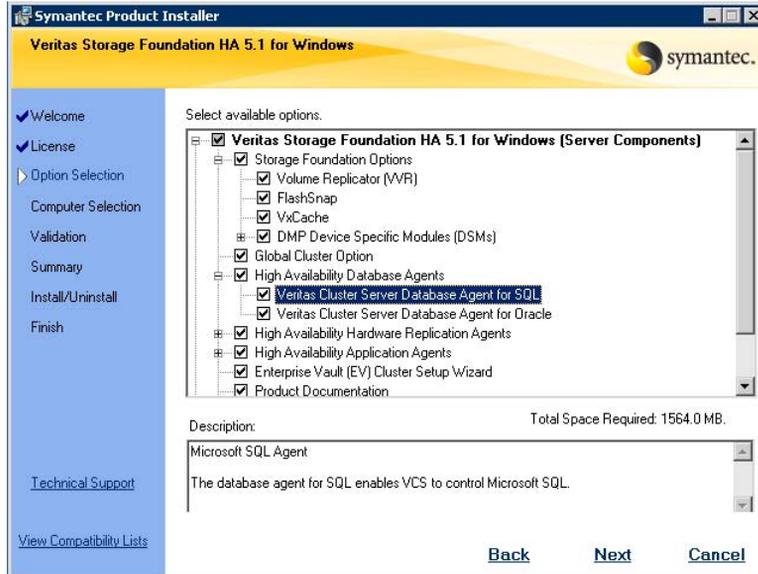
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

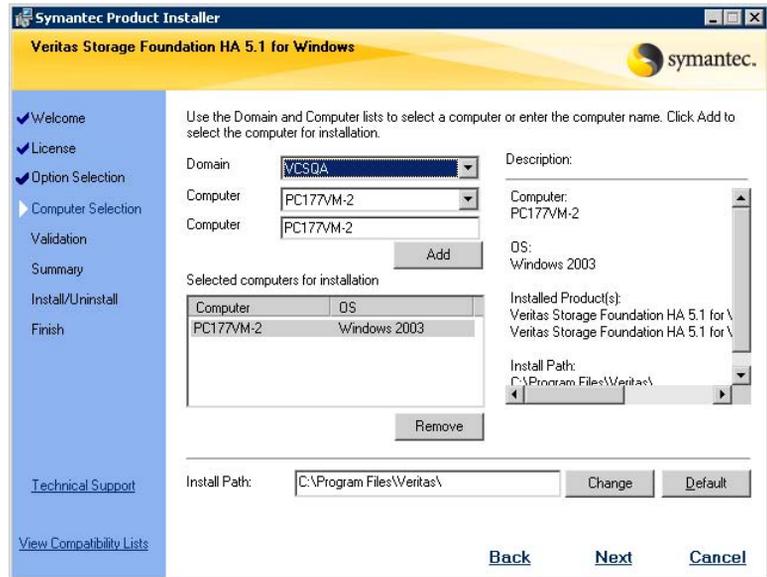
Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option Required for a disaster recovery configuration only.

Veritas Volume Replicator For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

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

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring cluster disk groups and volumes for SQL Server 2005

You must create a cluster disk group and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [About cluster disk groups and volumes](#)
- [Prerequisites for configuring cluster disk groups and volumes](#)
- [Sample disk group and volume configuration](#)
- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)
- [Creating a cluster disk group](#)
- [Creating volumes](#)

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: Create a cluster disk group and volumes for the SQL instance on only one node of a cluster. You make the volumes accessible 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes for the SQL instance:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process:

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

For additional disk group and volume requirements, see the following topics:

- [Considerations for disk groups and volumes for multiple instances](#)
- [Considerations for volumes for a disaster recovery configuration](#)

Considerations for disk groups and volumes for multiple instances

For an active-active configuration or other cases where you are setting up multiple SQL instances in the cluster, you create a separate set of cluster disk groups and volumes for each instance.

For example, if you have a Billing instance and a Payroll instance, you could create the following disk groups and volumes.

For the Billing instance, create the following:

- BILLING_DG: a cluster disk group for the volumes related to the Billing instance
- BILLING_DATA_FILES: volume for the SQL Server system data files
- BILLING_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Billing instance
- BILLING_DB1_VOL: volume for the user database files
- BILLING_DB1_LOG: volume for the user database log files

For the Payroll Instance, create the following:

- PAYROLL_DG: a cluster disk group for the volumes related to the Payroll instance
- PAYROLL_DATA_FILES: volume for the SQL Server system data files
- PAYROLL_REGREP_VOL: volume for the list of registry keys replicated among cluster nodes for the Payroll instance
- PAYROLL_DB1_VOL: volume for the user database files
- PAYROLL_DB1_LOG: volume for the user database log files

You can choose either of the following:

- Set up disk groups and volumes for all instances at one time.
- Set up disk groups and volumes for the current instance only and complete all configuration steps for this instance. Then return to this step for the next instance.

Considerations for volumes for a disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- A disaster recovery configuration with VVR requires a Storage Replicator Log (SRL) volume 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. For more about VVR planning, see the *Veritas Volume Replicator Administrator's Guide*.

- Symantec recommends that for replication considerations, you create a separate volume for tempdb, for example, INST1_TEMPDB, within the system database disk group. When you later configure replication for disaster recovery, you replicate that disk group but exclude the tempdb volume from the replication.

It would waste bandwidth to replicate tempdb because the data is transitory and is not needed for DR site recovery.

You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume. See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 272.

- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Creating a cluster disk group

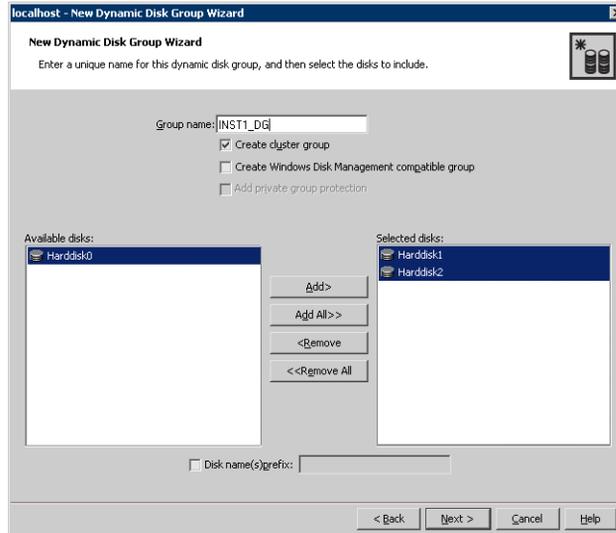
Use the Veritas Enterprise Administrator (VEA) to create a cluster disk group on the first node where the SQL instance is being installed. Repeat the procedure if you want to create additional disk groups.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

- 8 Click **Finish** to create the new disk group.

Creating volumes

This procedure will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure to create additional volumes. For example, create the following volumes on the INST1-DG disk group:

- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL**: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

You could create the following volumes on the same disk group as the system volumes or on a separate disk group, depending on how you prefer to configure your storage:

- **INST1_DB1_VOL**: contains the user database files
- **INST1_DB1_LOG**: contains the user database log files

Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.

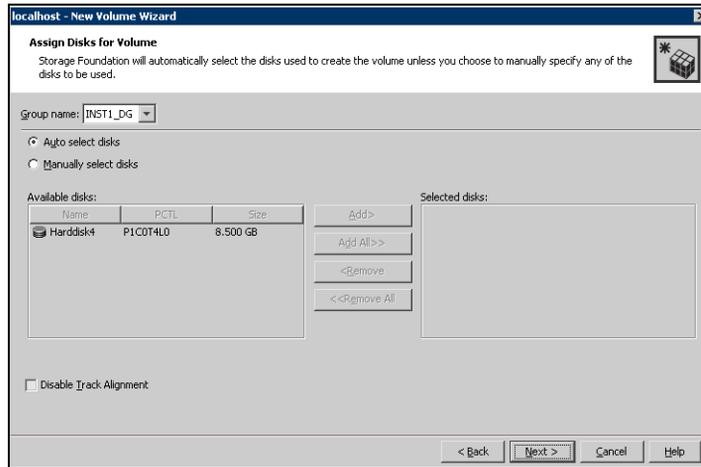
See “[Considerations for volumes for a disaster recovery configuration](#)” on page 232

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 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.

- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

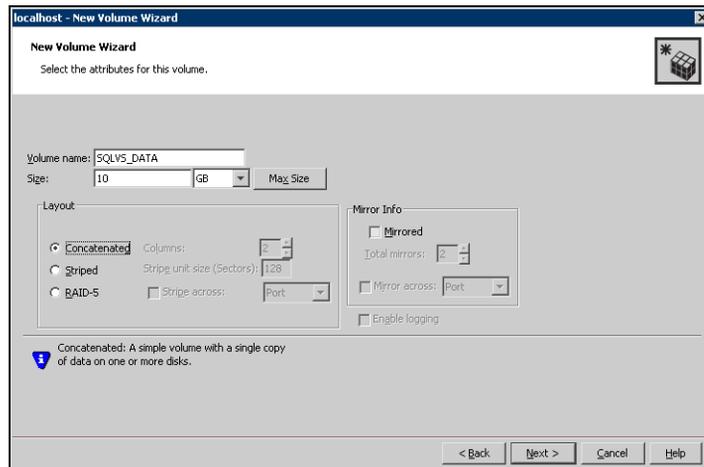


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

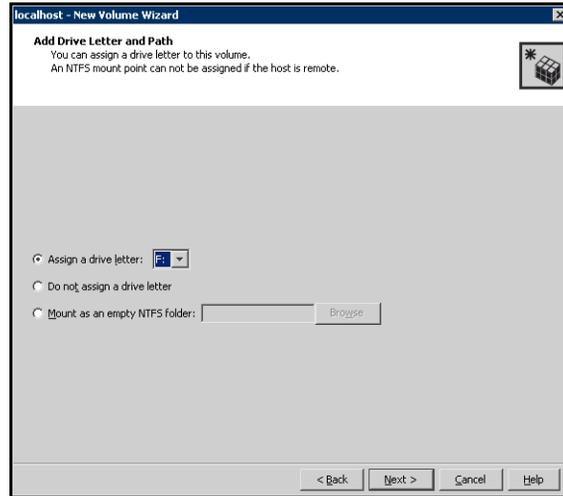
- 8 Click **Next**.

9 Specify the parameters of the volume.



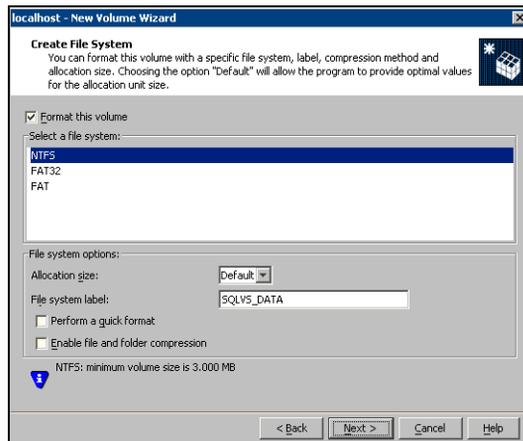
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Configuring the cluster

The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also provides the option to configure the ClusterService group, which can contain resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 205.

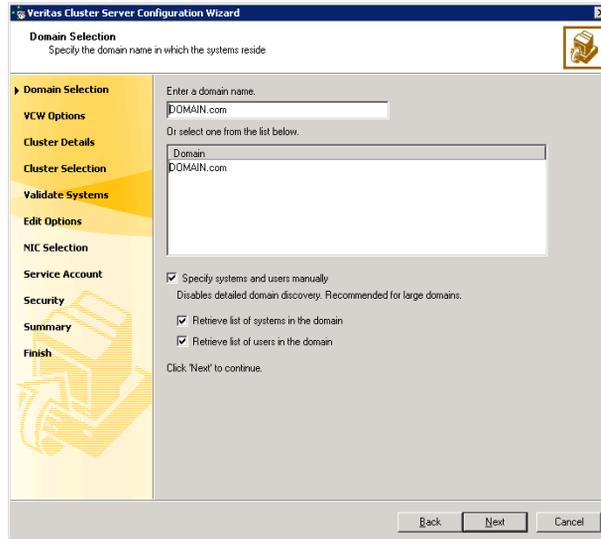
Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or removing or modifying cluster configurations.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

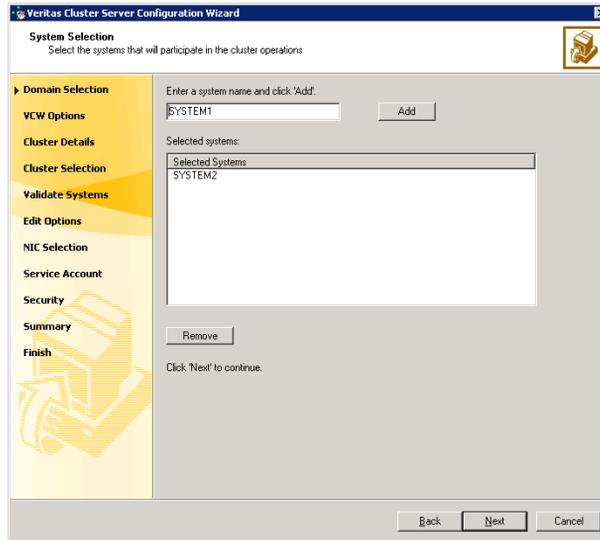
Proceed to [step 8](#) on page 242.

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

- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

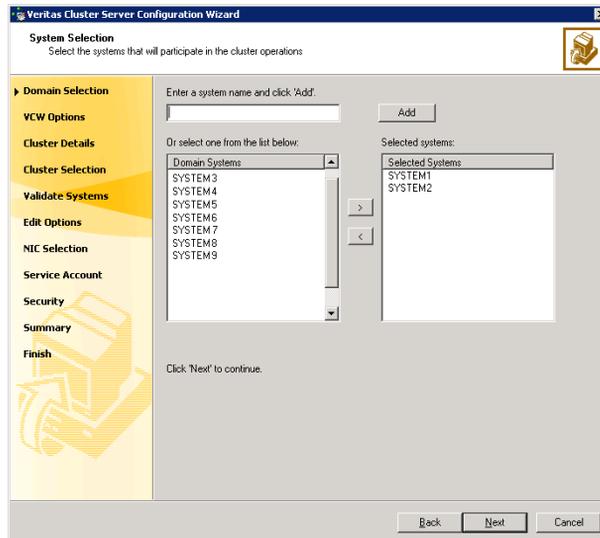
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 241. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 242.

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



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

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

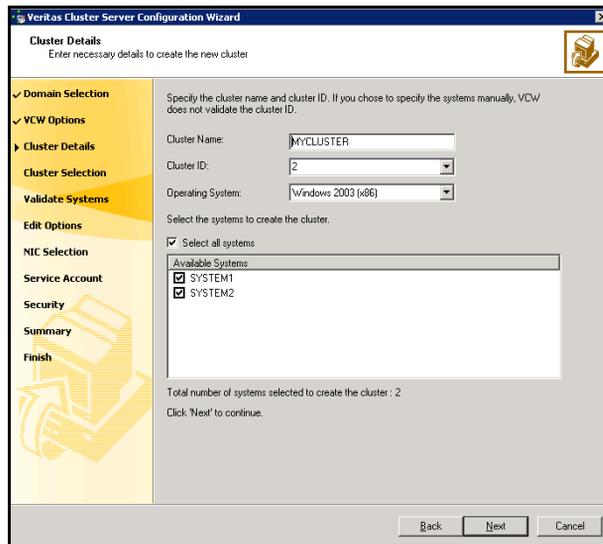
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 240 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

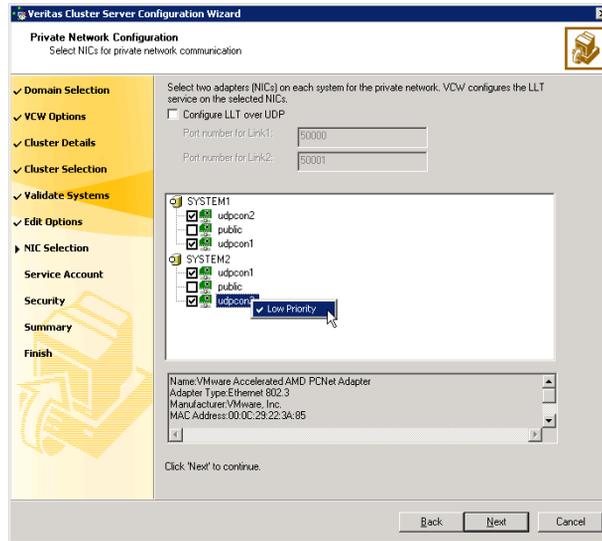
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 242, proceed to the next step. Otherwise, proceed to [step 12](#) on page 246.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

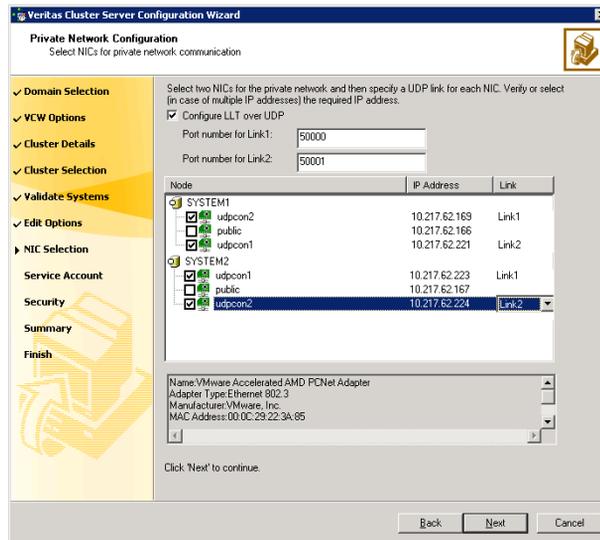
- To configure the VCS private network over Ethernet



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

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

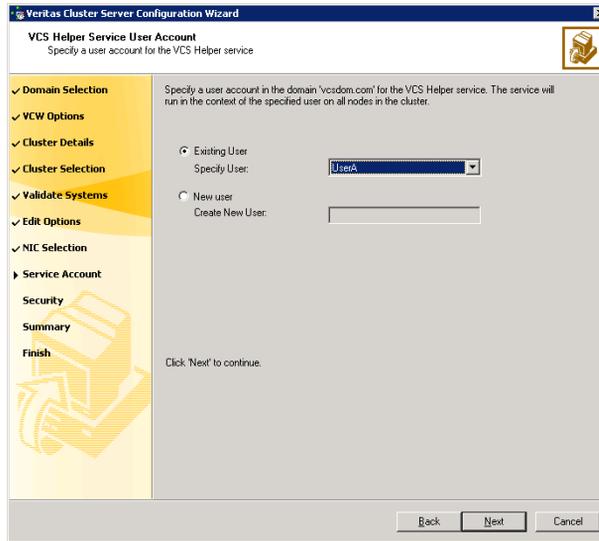
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 240, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

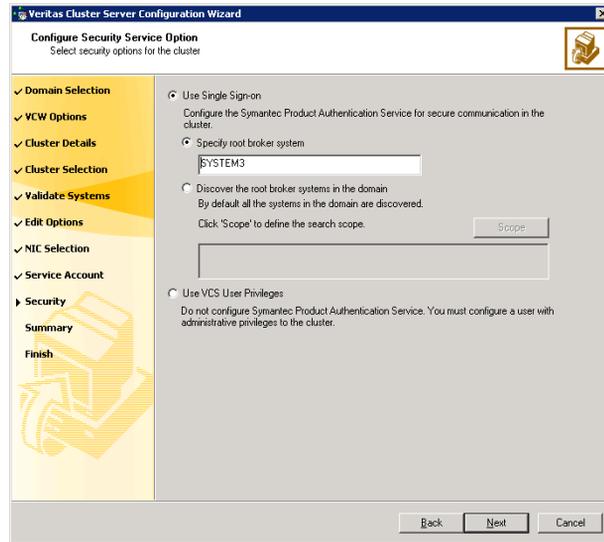
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 6-10](#) on page 248 contains some more examples of search criteria.

Table 6-10 Search criteria examples

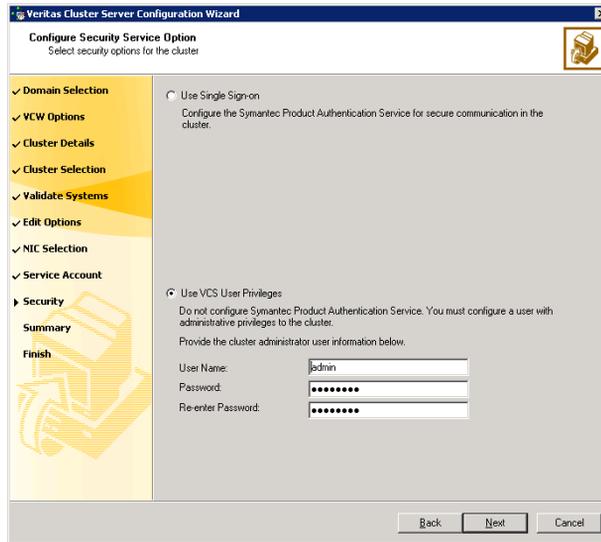
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for

the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

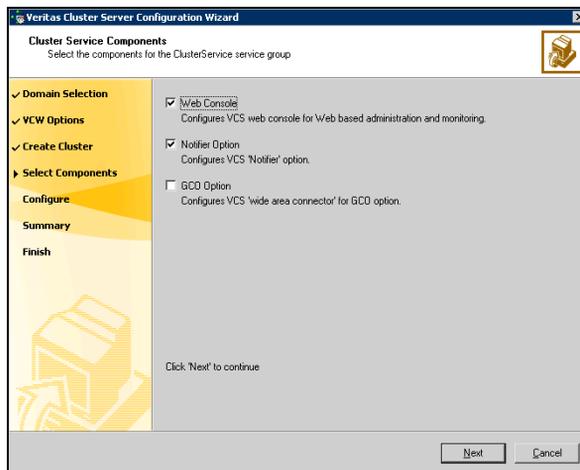
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

See “[Configuring Web console](#)” on page 251.

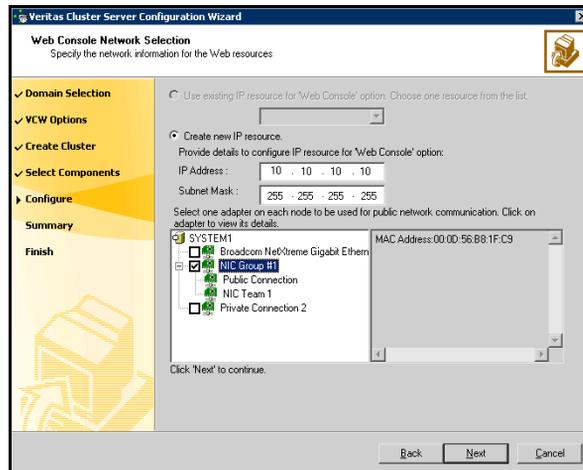
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
 See “[Configuring notification](#)” on page 252.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to:

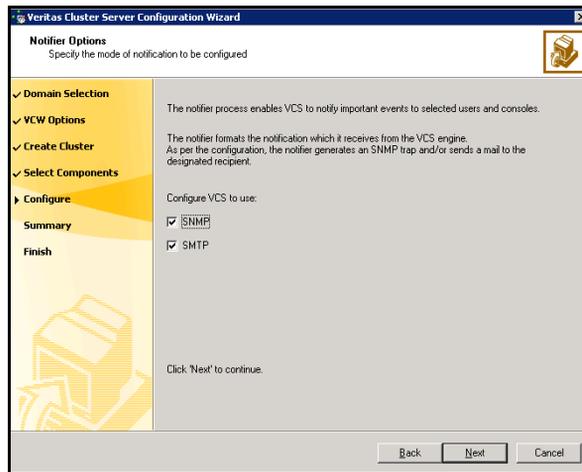
“Configuring notification” on page 252.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

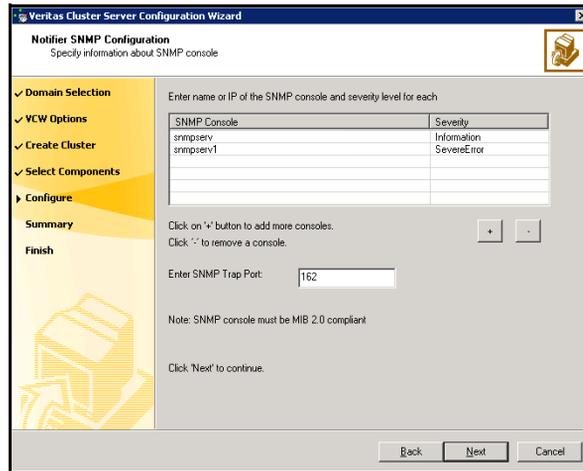
To configure notification

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



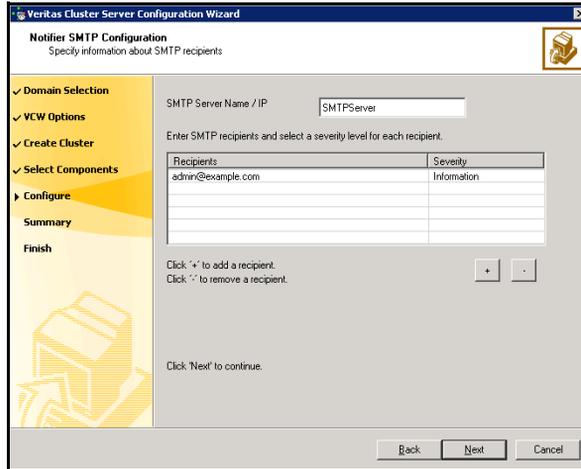
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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



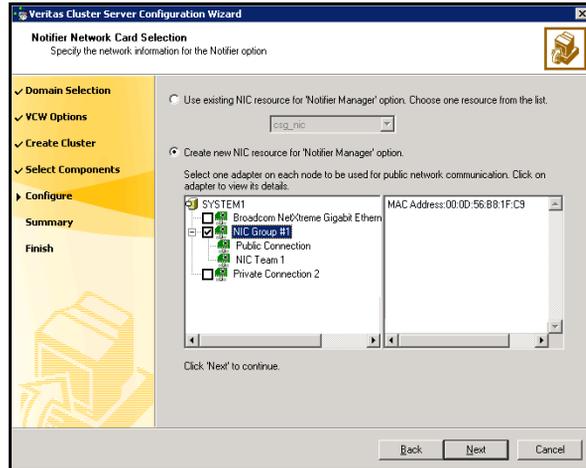
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



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

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



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

About installing multiple instances

If you are installing multiple instances of SQL Server on the same system, as in an active-active cluster configuration, some additional requirements apply. Procedures include these requirements where necessary. The following summary is provided for your review to assist you in planning the installation:

- Symantec recommends that you follow all steps for installing and setting up high availability for the first instance before you begin installing the next instance.
- Multiple instances of SQL Server must be installed in the same order on every node in the cluster. For example, if you install INST1 on SYSTEM1 and then on SYSTEM2, then install INST2 on SYSTEM1 and then on SYSTEM2.
- Assign a unique name to each instance. When installing SQL Server on additional nodes for the same instance, use the same instance name.
- Set a unique virtual server name as the internal name for each clustered instance.
- Assign a unique port number for each instance.

Installing and configuring SQL Server 2005 on the first node

In preparation for installing Microsoft SQL Server 2005, ensure that the cluster disk group is imported to the first node for this SQL instance and the volumes are mounted. See [“Importing the cluster disk group”](#) on page 262 and [“Adding drive letters to mount the volumes”](#) on page 263.

Complete the following procedures to install and configure this instance of Microsoft SQL Server 2005:

- [Installing Microsoft SQL Server 2005 on the first node](#)
- [Setting the startup mode of the SQL Server 2005 services](#)

Installing Microsoft SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

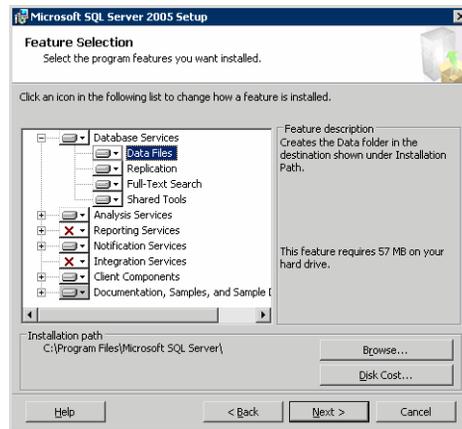
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

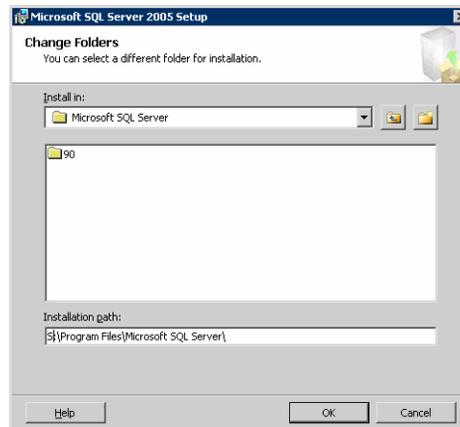
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.



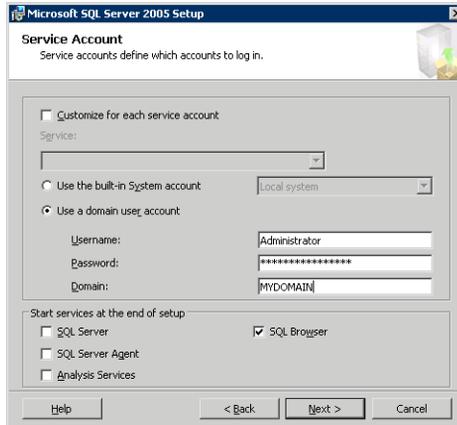
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 257, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.

Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.

- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

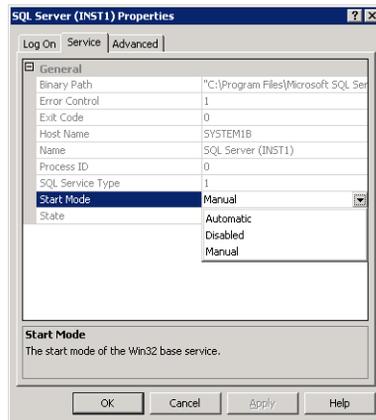
- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic. Refer to the steps described in the procedure that follows.

Setting the startup mode of the SQL Server 2005 services

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Complete the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:

- [“Stopping the SQL Server 2005 service”](#) on page 261
- [“Deporting the cluster disk group”](#) on page 262
- [“Importing the cluster disk group”](#) on page 262
- [“Adding drive letters to mount the volumes”](#) on page 263
- [“Renaming shared SQL Server 2005 files”](#) on page 264

Stopping the SQL Server 2005 service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

To install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name where the disk group is current imported, expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (for example, INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (for example, INST1_DG) to the next node in the cluster.

To import a cluster disk group

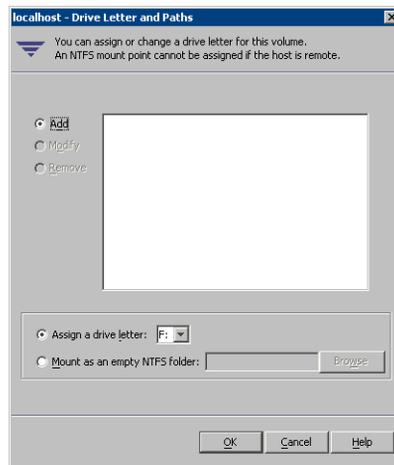
- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (for example, INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this instance:

- [“Installing Microsoft SQL Server on the second node”](#) on page 264
- [“Setting the internal name of the clustered instance”](#) on page 268

Installing Microsoft SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the additional node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 262 and [“Adding drive letters to mount the volumes”](#) on page 263.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Multiple instances of SQL Server 2005 must be installed in the same order on every node of the cluster.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

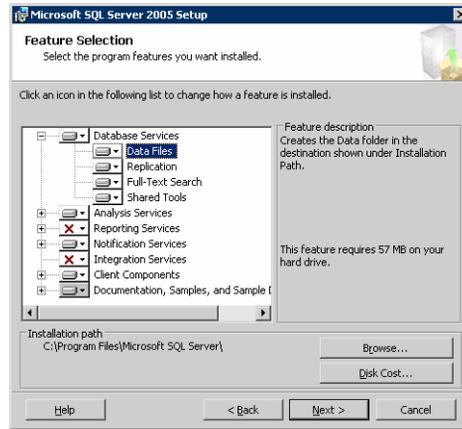
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

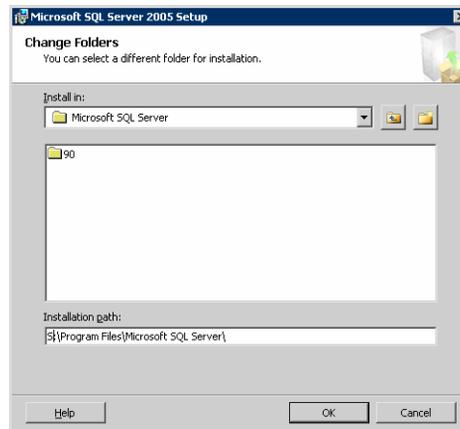
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.



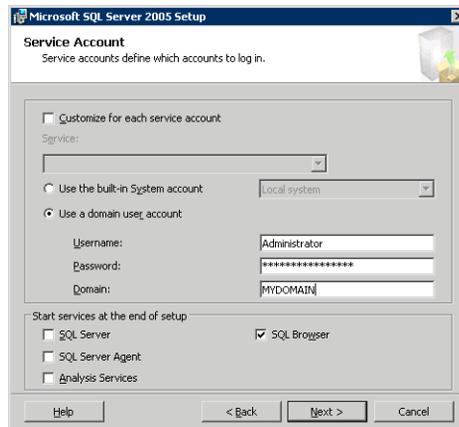
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 265, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.

Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.

- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.
 See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.
- Refer to “[Setting the startup mode of the SQL Server 2005 services](#)” on page 260 for the procedure.

Repeat the procedures described in “[Preparing to install SQL Server 2005 on the second node](#)” on page 261 and “[Installing and configuring SQL Server 2005 on the second node](#)” on page 264 on any additional nodes for the same SQL instance.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, `INST1-VS\INST1`).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

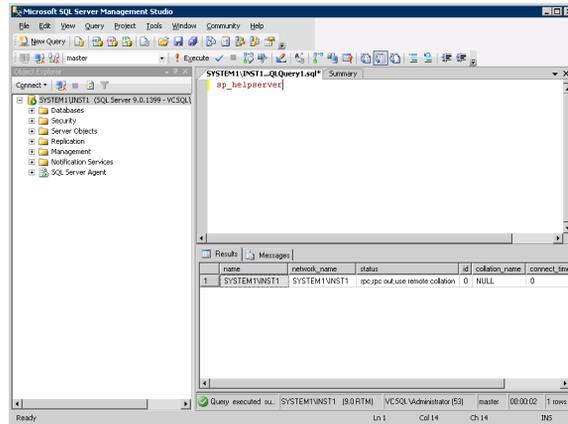
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, *SYSTEM1\INST1*. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
sp_dropserver "System_Name\Instance_Name"
where *System_Name\Instance_Name* is the name noted in [step 3](#) on page 269.
For example, for a named instance:
sp_dropserver "SYSTEM1\INST1"
For example, for a default instance:
sp_dropserver "SYSTEM1"
 - Press F5.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database.

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See “[Creating volumes](#)” on page 235.

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Connect to the SQL Server virtual server instance for which you want to create the database (for example, `INST1-VS\INST1`).
- 3 Expand the icon associated with your server.
- 4 Right-click on **Databases** and select **New Database**.
- 5 In the New Database page, enter a name for the new database.
- 6 Click the browse button (...) in the **Path** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 7 Select and edit other file properties as desired.
- 8 Click the browse button (...) in the **Path** column for the **Transaction Log** row and browse to the location of the volume you want to create for the transaction log, and click **OK**.
- 9 Add more data files if required and configure whatever options are required for your database.
- 10 Click **OK**.

- 11 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See “[Completing configuration steps in SQL Server](#)” on page 272.
- 12 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See “[Modifying a SQL 2005 service group to add VMDg and MountV resources](#)” on page 286.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume within the system database disk group.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 272.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See “[Assigning ports for multiple SQL Server instances](#)” on page 273.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating volumes](#)” on page 235.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Assigning ports for multiple SQL Server instances

If you are running multiple SQL Server instances, you must assign a different port to each named instance. You can assign static or dynamic ports.

Refer to the Microsoft Knowledge Base for the instructions on assigning ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

If you wish to change the port after configuring the SQL service group, you must perform the steps in the following order:

- Bring the SQL service group online or partially online (up to the registry replication resource) on a cluster node.
- On the node on which the SQL service group is online or partially online, change the port assigned to the SQL instance. Refer to the instructions mentioned in the Microsoft Knowledge Base article specified earlier.
- Take the SQL service group offline on the node, and then bring it online again. The configuration changes will be replicated to the remaining cluster nodes.

Configuring the VCS SQL Server 2005 service group

A VCS SQL Server service group is used to bring a SQL Server 2005 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Read the following topics:

- [Service group requirements for active-active configurations](#)
- [Prerequisites for configuring the service group](#)
- [Creating the SQL Server 2005 service group](#)

Service group requirements for active-active configurations

Note the following requirements for active-active configurations:

- For an active-active configuration, you must create a separate service group for each instance.
- Each service group that you create must have a unique service group name and virtual IP address.
- For an active-active configuration, when you specify the priority order of systems, reverse the order for each service group so that the active system and failover system are opposite for each instance. For example, if you have two instances and two systems, you would set the priority order as follows:

INSTANCE 1	Priority order: SYSTEM1 SYSTEM2
INSTANCE2	Priority order: SYSTEM2 SYSTEM1

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- Verify that SFW HA, along with the VCS database agent for SQL Server 2005, is installed on all cluster nodes. See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 223.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 239.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group. See [“Importing the cluster disk group”](#) on page 262 and [“Adding drive letters to mount the volumes”](#) on page 263.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. You specify the virtual server name when configuring the service group. See [“Setting the internal name of the clustered instance”](#) on page 268.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.

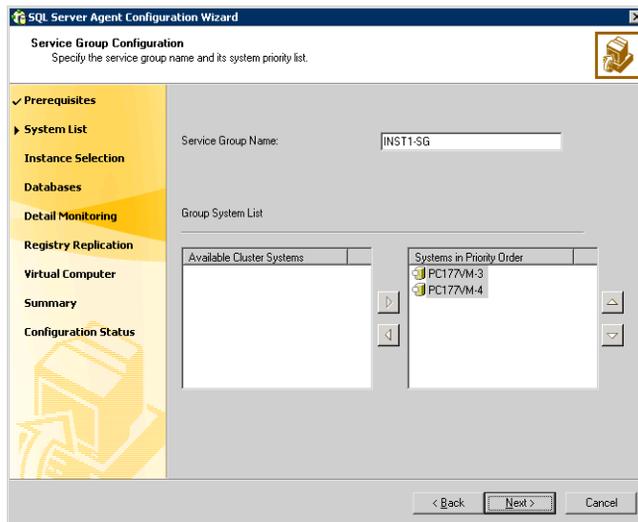
- Assign a unique virtual IP address for the SQL Server 2005 instance. You specify this IP address when configuring the service group.
- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes. A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2005\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2005 Server service for the SQL instance. See [“Stopping the SQL Server 2005 service”](#) on page 261.

Creating the SQL Server 2005 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:

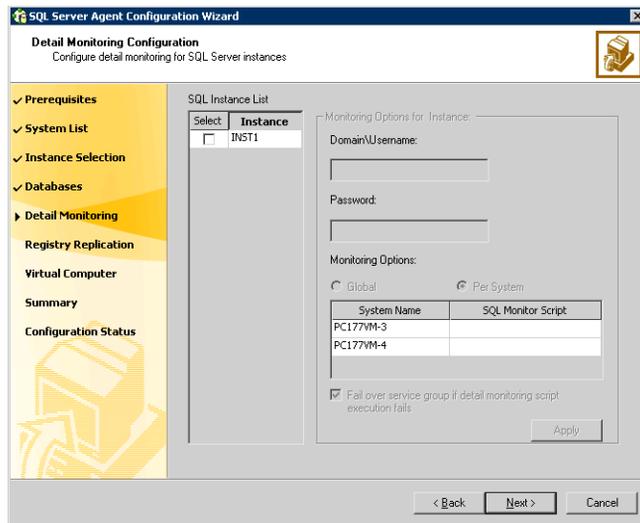


- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the

systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.

For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

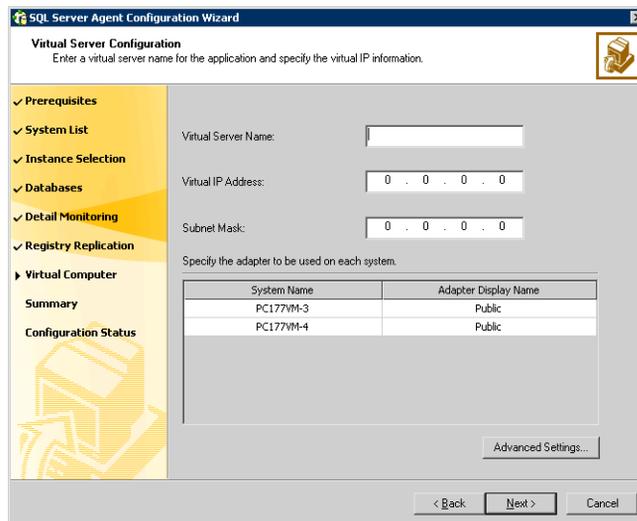
- Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
 Click **Next**.
 - 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
 - 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 277 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script

on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.

- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 277 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.

- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
 The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.

- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.

15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.

- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
 You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

If you have created a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

See “[Modifying a SQL 2005 service group to add VMDg and MountV resources](#)” on page 286.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for high availability](#)” on page 355.

Verifying the SQL Server 2005 cluster configuration

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

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

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in step 1.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.

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

Determining additional steps needed

Determine the next step as follows:

- If you have no more instances of SQL Server that you plan to install and configure for high availability, your high availability configuration is complete.
- If you have additional instances of SQL Server to install and configure in the cluster, start the installation and configuration sequence again for the next instance.
 - If you are using the Solutions Configuration Center:
See [“Following the workflow in the Solutions Configuration Center”](#) on page 219.
 - If you are using the printed document as a guide:
See [“Tasks for a new HA installation of SQL Server 2005”](#) on page 200.
- If you are planning to set up a secondary site for disaster recovery:
See [Chapter 17, “Deploying disaster recovery: New SQL Server 2005 installation”](#) on page 881.
- If you want to configure a Cluster Management Console connection:
See [“Configuring the Cluster Management Console connection”](#) on page 282.
- If you need to modify the SQL Server service group:
See [“Modifying a SQL 2005 service group to add VMDg and MountV resources”](#) on page 286.
- If you need to make other changes to the cluster:
See *Veritas Cluster Server Administrator’s Guide*.

Configuring the Cluster Management Console connection

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

CMC comprises of the following components:

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

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

Prerequisites for installing the cluster connector

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

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

Installing the cluster connector on Windows clusters

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

To install the cluster connector on a Windows cluster

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

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

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

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

Configuring the cluster connector

Perform the following steps to configure the cluster connector.

To configure the cluster connector

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

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

This indicates the configuration is successful.

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

Configuring the cluster connector using the management server console

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

To upgrade the cluster connector on discovered clusters

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

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

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

Modifying a SQL 2005 service group to add VMDg and MountV resources

If you create a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.

- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as **CRITICAL**.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for high availability](#)” on page 355.

Deploying SFW HA for high availability: Standalone SQL 2005 Servers

This chapter includes the following topics:

- [“Tasks for converting a standalone SQL 2005 Server for high availability”](#) on page 290
- [“Reviewing the requirements”](#) on page 293
- [“Reviewing the configuration”](#) on page 298
- [“Configuring the storage hardware and network”](#) on page 300
- [“Preparing the standalone SQL Server”](#) on page 302
- [“Installing Veritas Storage Foundation for Windows”](#) on page 304
- [“Configuring cluster disk groups and volumes for SQL Server 2005”](#) on page 310
- [“Configuring the cluster”](#) on page 320
- [“Verifying that SQL Server 2005 databases and logs are moved to shared storage”](#) on page 336
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 336
- [“Installing and configuring SQL Server 2005 on the second node”](#) on page 339

- “[Configuring the VCS SQL Server 2005 service group](#)” on page 346
- “[Verifying the SQL Server 2005 cluster configuration](#)” on page 351
- “[Verifying the SQL Server 2005 cluster configuration](#)” on page 351
- “[Adding a new SQL Server user-defined database](#)” on page 353
- “[Additional instructions for disaster recovery](#)” on page 353

Tasks for converting a standalone SQL 2005 Server for high availability

This chapter describes the procedure to convert a standalone SQL 2005 Server into a “clustered” SQL Server in a new Veritas Storage Foundation HA environment. This environment involves an active-passive configuration with one to one failover capabilities.

To plan a new SQL Server 2005 deployment, or to review additional considerations for an active-active configuration, refer to [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#) on page 199.

Note: In addition to the information contained in this chapter, the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach are required. Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Note: Some installation and configuration options are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery using Veritas Volume Replicator (VVR).

Symantec recommends using the Solutions Configuration Center as a guide for configuring high availability for SQL Server 2005.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

To configure MSDTC service groups, see [Chapter 8, “Configuring an MSDTC service group for high availability”](#) on page 355.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 7-1 Tasks for converting a standalone SQL 2005 server for high availability

Objective	Tasks
“ Reviewing the requirements ” on page 293	Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 298	<ul style="list-style-type: none"> ■ Understanding active-passive configuration ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 300	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Preparing the standalone SQL Server ” on page 302	<ul style="list-style-type: none"> ■ Backing up existing data ■ Setting SQL Server services to manual start
“ Installing Veritas Storage Foundation for Windows ” on page 304	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation HA for Windows (automatic installation) ■ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server
“ Configuring cluster disk groups and volumes for SQL Server 2005 ” on page 310	<ul style="list-style-type: none"> ■ Planning the storage layout ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ For a new shared storage configuration, creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator

Table 7-1 Tasks for converting a standalone SQL 2005 server for high availability (Continued)

Objective	Tasks
“Configuring the cluster” on page 320	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“Verifying that SQL Server 2005 databases and logs are moved to shared storage” on page 336	<ul style="list-style-type: none"> ■ Stopping SQL Server service ■ Modifying data file and user database locations
“Preparing to install SQL Server 2005 on the second node” on page 336	<ul style="list-style-type: none"> ■ Ensuring that the cluster disk group is imported on the additional node ■ Optionally, renaming system data files
“Installing and configuring SQL Server 2005 on the second node” on page 339	<ul style="list-style-type: none"> ■ Installing SQL Server on additional nodes ■ Setting the internal name of the clustered instance
“Configuring the VCS SQL Server 2005 service group” on page 346	<p>Creating a SQL Server service group using the VCS SQL Configuration wizard</p>
“Adding a new SQL Server user-defined database” on page 353	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server ■ Adding resources for a user-defined database in VCS
“Verifying the SQL Server 2005 cluster configuration” on page 351	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes

Reviewing the requirements

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
|--|--|

Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)	■	Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required)	■	Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
	■	Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)
	■	Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)
Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)	■	Windows Server 2008 for 64-bit Itanium (IA64)
	■	Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See “[Best practices](#)” on page 296.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP

addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on `C:\WINDOWS` of one node, installations on all other nodes must be on `C:\WINDOWS`. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.

- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Disk space requirements

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

[Table 7-2](#) on page 297 estimates disk space requirements for SFW HA.

Table 7-2 Disk space requirements

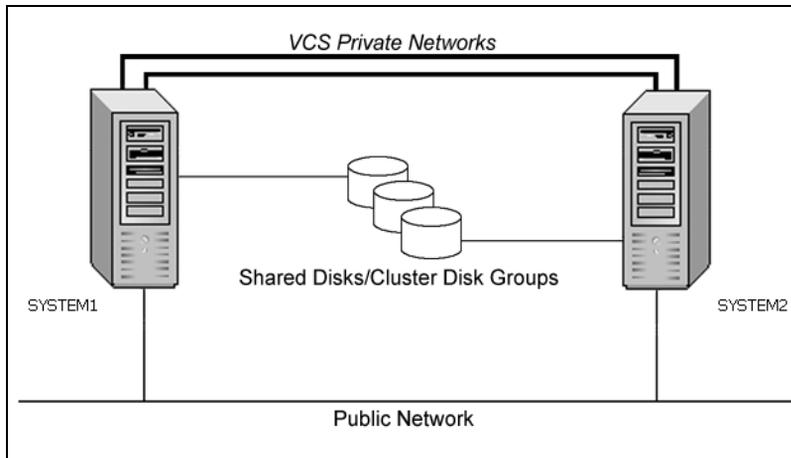
Installation options	Install directory/drive
SFW HA + all options + client components	1564 MB
SFW HA + all options	1197 MB
Client components	528 MB

Reviewing the configuration

This chapter describes the tasks needed to incorporate an existing standalone SQL Server into a high available environment in order to ensure that the mission critical SQL resource is always available.

This chapter describes the tasks necessary to create a virtual server in an Active-Passive SQL configuration. 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. At the end of this process, their environment will look like this:

Figure 7-1 Active-Passive configuration



The virtual SQL Server is online on SYSTEM1, serving client requests. The shared disks provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. From the user's perspective there will be a small delay as the backup node comes online, but the interruption in effective service is minimized.

Sample configuration

A sample setup is used through this guide to illustrate the installation and configuration tasks.

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the IP address should be the same on all nodes
- Cluster IP address: used by Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console

You should have these IP addresses available before you start deploying your environment. The following names describe the objects created and used during the installation and configuration:

Table 7-3 Standalone SQL Server 2005 configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL Server
INST1_SG	Microsoft SQL Server 2005 service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1	SQL Instance Name
INST1-VS	name of the SQL Virtual Server

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Preparing the standalone SQL Server

Complete the following tasks before you begin the process of installing Veritas Storage Foundation HA for Windows and creating a clustered environment:

- Backing up existing SQL data
- Setting SQL Server services to manual start

Backing up existing SQL data

Create a backup of the data on the existing standalone SQL Server.

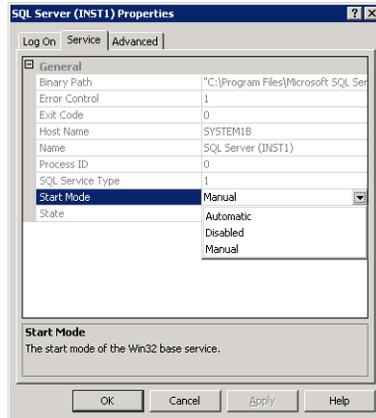
Setting SQL Server services to manual start

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).

- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Installing Veritas Storage Foundation for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 7-4](#) on page 304 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 7-4 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 305.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

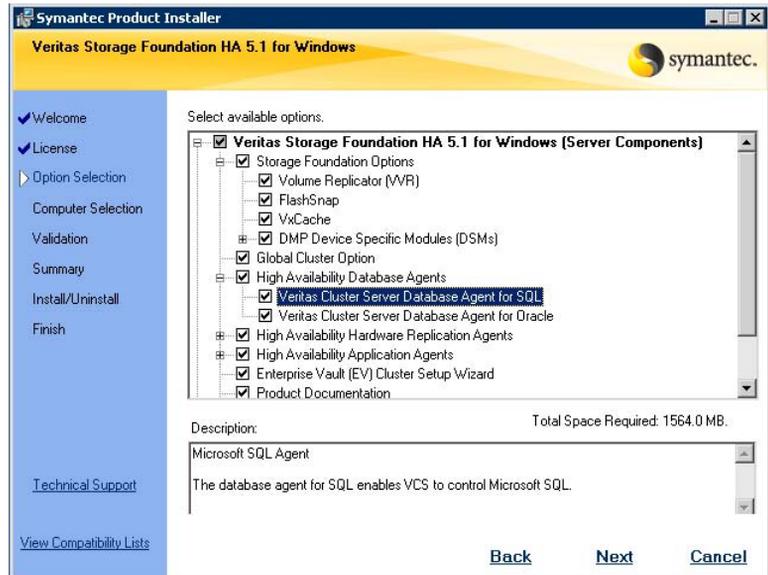
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

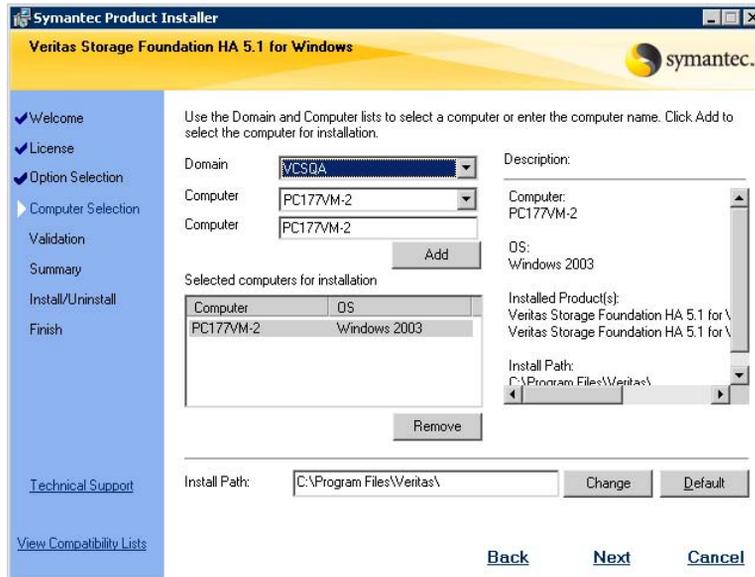
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



- | | |
|---|---|
| Veritas Cluster Server Database Agent for SQL | Required to configure high availability for SQL Server. |
| Client | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Veritas Volume Replicator | For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR. |
| High Availability Hardware Replication Agents | If you plan to use hardware replication, select the appropriate hardware replication agent. |

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Configuring cluster disk groups and volumes for SQL Server 2005

You create cluster disk groups and volumes to manage your SQL Server database and logs, as covered in the following topics:

- [“About cluster disk groups and volumes”](#) on page 310
- [“Prerequisites for configuring cluster disk groups and volumes”](#) on page 311
- [“Sample disk group and volume configuration”](#) on page 312
- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 312
- [“Considerations for volumes for a VVR disaster recovery configuration”](#) on page 313
- [“Creating a cluster disk group”](#) on page 314
- [“Creating volumes”](#) on page 316

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: Create a cluster disk group and volumes for the SQL instance on only one node of a cluster. You make the volumes accessible 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - *Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring cluster disk groups and volumes

Complete the following tasks before you create the cluster disk group and volumes for the SQL instance:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
See “[Sample disk group and volume configuration](#)” on page 312.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

You may be configuring new shared storage for the high availability environment, or the existing standalone SQL Server databases and logs may already be on shared storage. If the existing databases and logs are already on shared storage, read the following topic:

- “[Considerations for converting existing shared storage to cluster disk groups and volumes](#)” on page 312

For a disaster recovery configuration using Veritas Volume Replicator, read the following topic:

- “[Considerations for volumes for a VVR disaster recovery configuration](#)” on page 313

Sample disk group and volume configuration

You first create a cluster disk group (INST1_DG) on shared disks. You then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process:

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

Your configuration may require additional planning. See the following topics:

- [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 312
- [“Considerations for volumes for a VVR disaster recovery configuration”](#) on page 313

Considerations for converting existing shared storage to cluster disk groups and volumes

The databases and logs for your existing standalone SQL Server may already be on shared storage. In this case, when you create cluster disk groups, you specify the disks that contain the existing databases and logs.

Creating a disk group converts the disks from basic disks to dynamic disks. Partitions on the disks are automatically converted to volumes on the dynamic disks.

Therefore, if your existing disk layout contains databases and logs in the same partition, they become part of the same volume in the cluster disk group. If the disk contains multiple partitions, each containing a user database, each partition becomes a separate volume, but all will become part of the same cluster disk group. If this configuration does not meet your requirements, you may want to modify your disk layout before creating the cluster disk group.

For additional information on converting basic to dynamic disks, see *Veritas Storage Foundation Administrator's Guide*.

Symantec recommends creating a separate 100 MB RegRep volume that contains the list of registry keys that must be replicated among cluster systems for the SQL service. However, if no additional disks are available on the shared storage, you can specify an existing volume as the registry replication path during service group creation.

For a disaster recovery configuration using Veritas Volume Replicator, you will need to allow additional disk space for a Storage Replicator Log volume.

See "[Considerations for volumes for a VVR disaster recovery configuration](#)" on page 313.

Considerations for volumes for a VVR disaster recovery configuration

For a disaster recovery configuration using Veritas Volume Replicator, note the following:

- When you run the Disaster Recovery Wizard, a VVR Storage Replicator Log volume is created automatically for each disk group that contains volumes that are replicated. Ensure that you allow sufficient disk space for this volume. For more about VVR planning, see the *Veritas Volume Replicator, Administrator's Guide*.
- Symantec recommends that for replication considerations, you create a separate volume for tempdb, for example, INST1_TEMPDB, within the system database disk group. When you later configure replication for disaster recovery, you replicate that disk group but exclude the tempdb volume from the replication.

It would waste bandwidth to replicate tempdb because the data is transitory and is not needed for DR site recovery.

You can create the volume now and later, after the SQL installation is complete and before configuring replication, move tempdb to the volume. See "[Moving the tempdb database if using VVR for disaster recovery](#)" on page 272.

- VVR does not support the following types of volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Creating a cluster disk group

Use the Veritas Enterprise Administrator (VEA) to create a cluster disk group on the existing standalone SQL Server system. Repeat the procedure if you want to create additional disk groups.

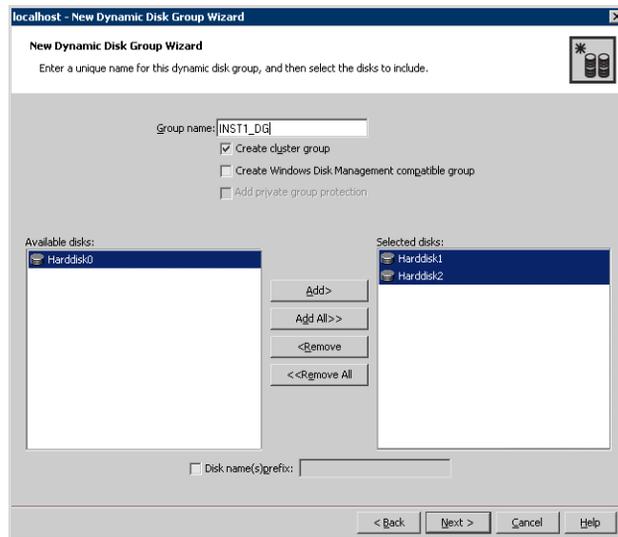
Ensure that you have reviewed all prerequisites and considerations before creating the cluster disk group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This procedure will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure to create additional volumes. For example, create the following volumes on the INST1-DG disk group:

- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL**: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

You could create the following volumes on the same disk group as the system volumes or on a separate disk group, depending on how you prefer to configure your storage:

- **INST1_DB1_VOL**: contains the user database files
- **INST1_DB1_LOG**: contains the user database log files

Additional considerations apply to configuring volumes for a disaster recovery configuration using VVR.

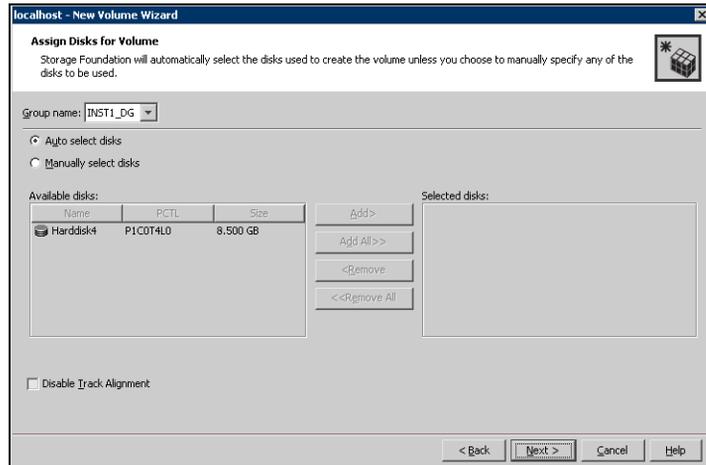
See “[Considerations for volumes for a VVR disaster recovery configuration](#)” on page 313.

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

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

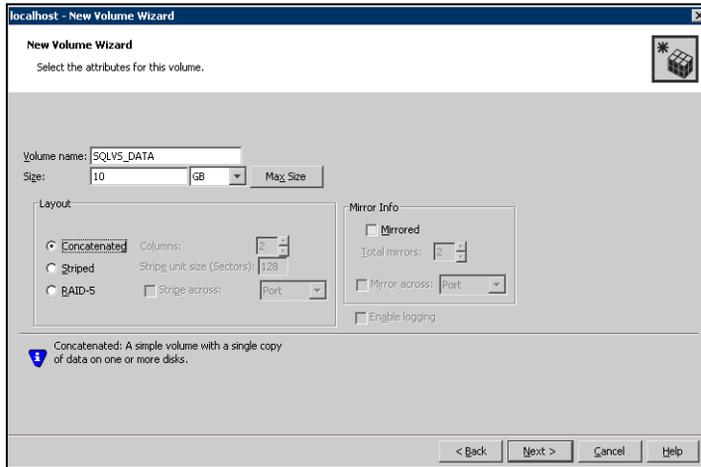


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

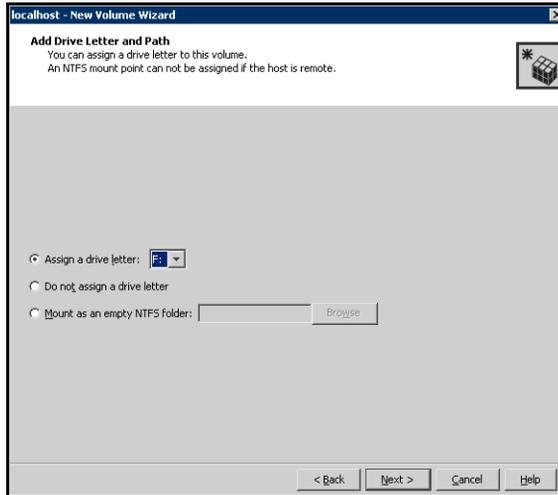
- 8 Click **Next**.

9 Specify the parameters of the volume.

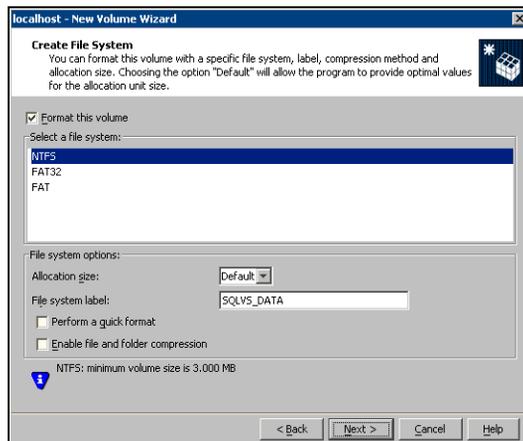


- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the Default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
- Click **Next**.

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

14 Repeat these steps to create additional volumes.

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

Configuring the cluster

You configure the cluster to contain the existing SQL Server 2005 system and the system that will become the failover node. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also provides the option to configure the ClusterService group, which can contain resources for Cluster Management Console (Single Cluster Mode) also referred to as Web Console, notification, and global clusters.

Complete the following tasks before creating a cluster:

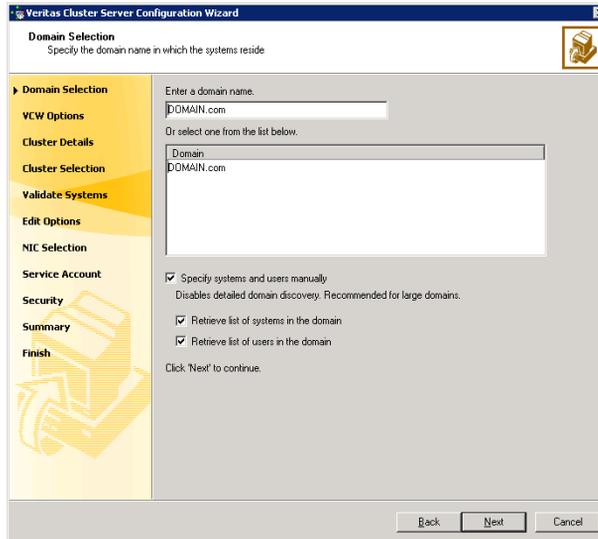
- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See “[Reviewing the requirements](#)” on page 293.

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS, including instructions on adding cluster nodes or removing or modifying cluster configurations.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



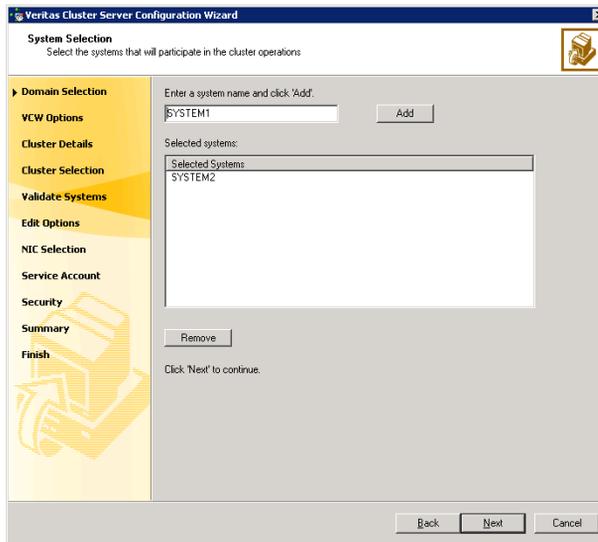
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
 Proceed to [step 8](#) on page 323.

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

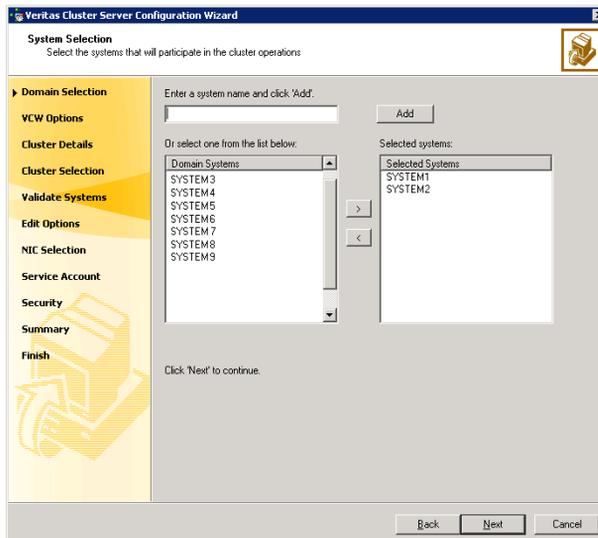
- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
 If you chose to retrieve the list of systems, proceed to [step 6](#) on page 322. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 323.

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



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

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

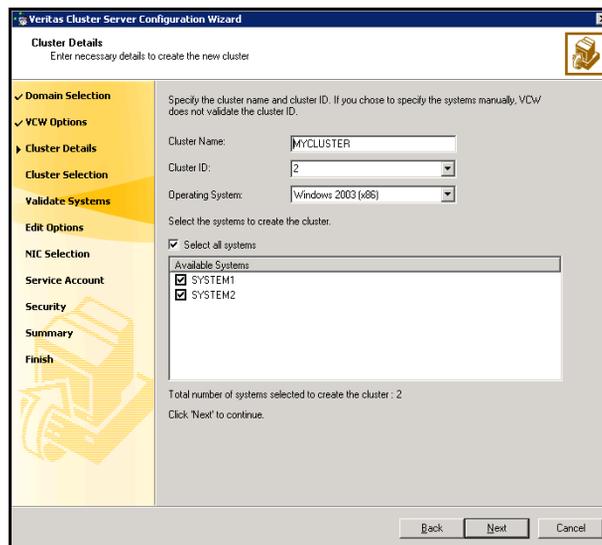
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 321 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

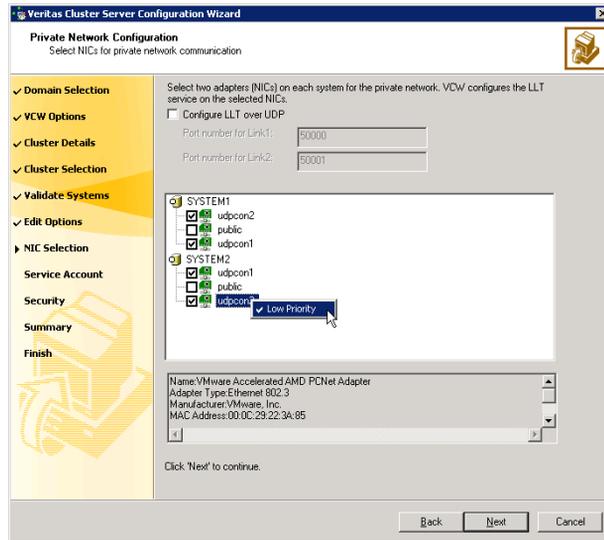
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 323, proceed to the next step. Otherwise, proceed to [step 12](#) on page 327.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

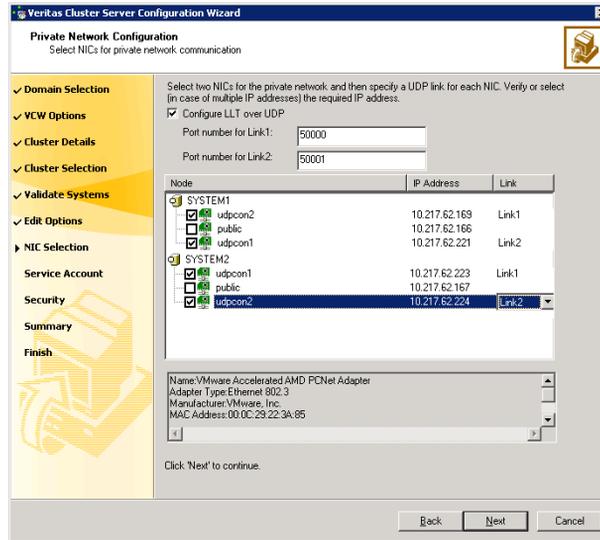
- To configure the VCS private network over Ethernet



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

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

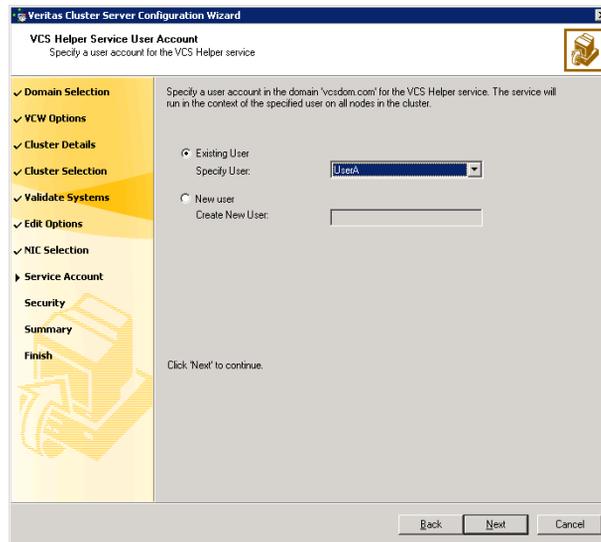
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



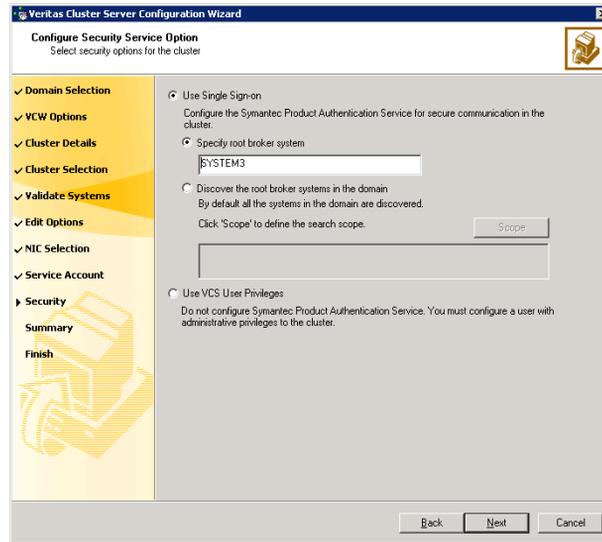
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 321, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
 Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities.

These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 7-5](#) on page 329 contains some more examples of search criteria.

Table 7-5 Search criteria examples

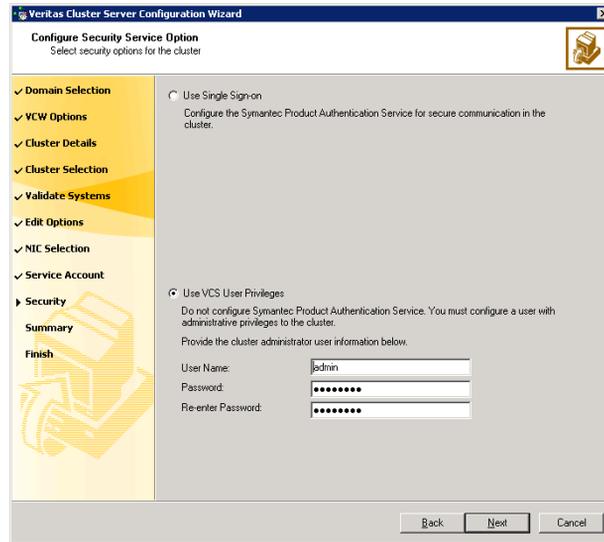
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.

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

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

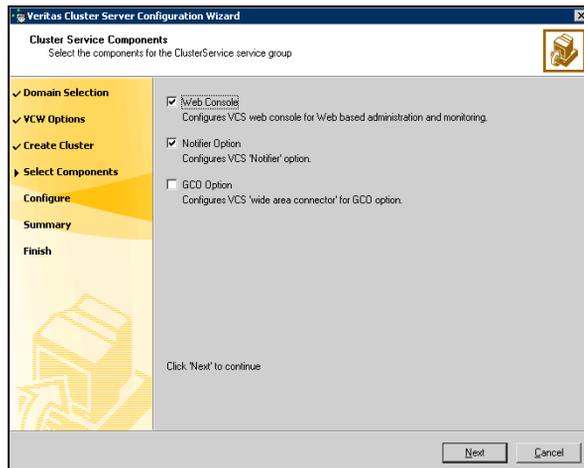
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 332.

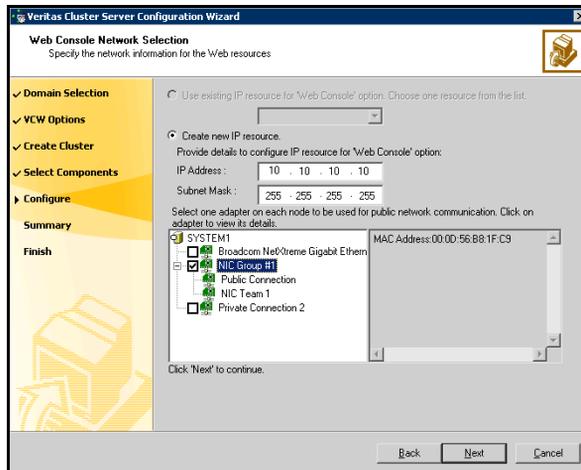
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See “[Configuring notification](#)” on page 333.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



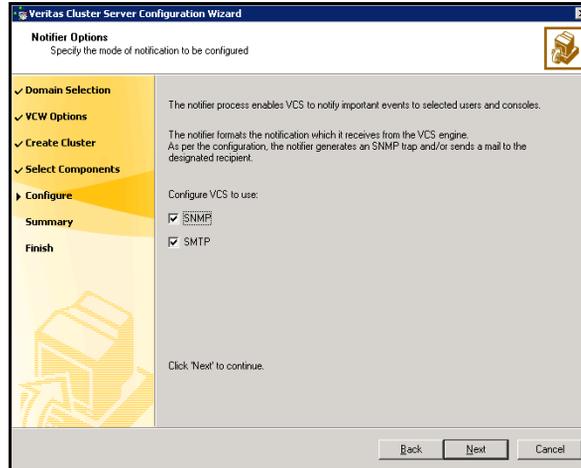
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 333.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

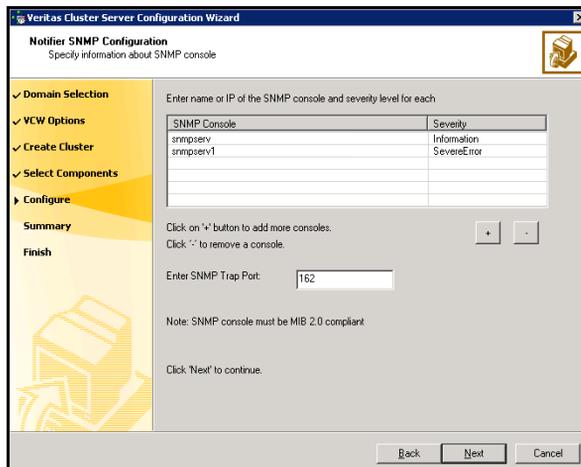
To configure notification

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

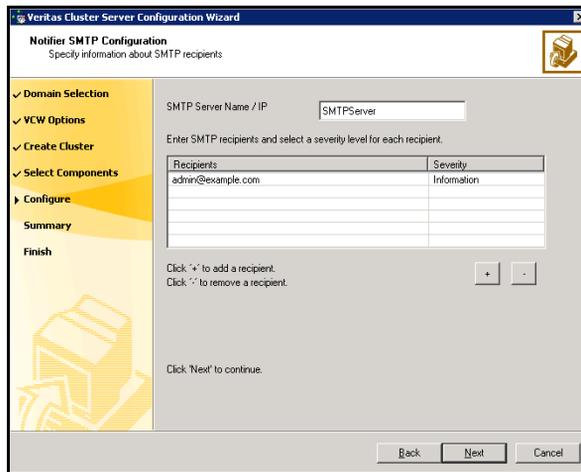


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

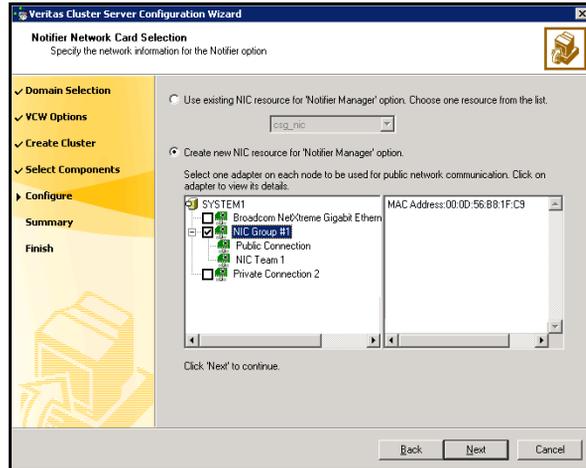


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Verifying that SQL Server 2005 databases and logs are moved to shared storage

Verify the location of all SQL Server databases and logs for the existing standalone server. If they are located on local storage, move them from the local drive to the appropriate cluster disk groups and volumes on shared storage to ensure proper failover operations in the cluster.

Complete the following tasks to move the databases:

- 1 Stop the SQL Server service.
- 2 Verify that you have backed up your existing data.
- 3 Ensure that the dynamic disk group is imported on the node where the original files are located on the local drives and mount the volumes. Refer to [“Deporting the cluster disk group”](#) on page 262, [“Importing the cluster disk group”](#) on page 262, and [“Adding drive letters to mount the volumes”](#) on page 263 for instructions.
- 4 Modify the SQL Server 2005 data file and user database locations. Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.
<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>
- 5 Restart SQL Server 2005.

Preparing to install SQL Server 2005 on the second node

Ensure that the cluster disk group is imported on the node where you install SQL Server 2000.

In addition, you may want to rename the system data files before installation, so that they can be restored if anything goes wrong with the installation.

See the following topics:

- [“Deporting the cluster disk group”](#) on page 337
- [“Importing the cluster disk group”](#) on page 337
- [“Adding drive letters to mount the volumes”](#) on page 338
- [“Renaming shared SQL Server 2005 files”](#) on page 339

Deporting the cluster disk group

To install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, you use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name where the disk group is current imported, expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (for example, INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (for example, INST1_DG) to the next node in the cluster.

To import a cluster disk group

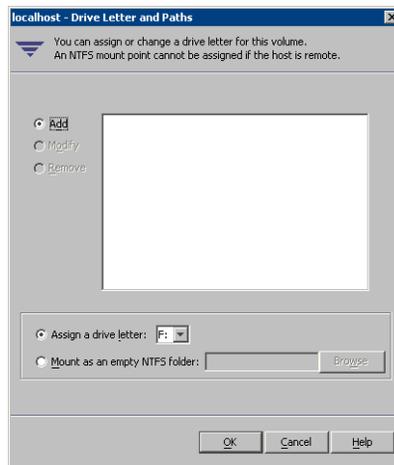
- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (for example, INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2005 on the second node

Install Microsoft SQL Server on each additional node that will be included in the SQL Server service group. Make sure to install the SQL Server data files to shared storage.

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2005 on the second node”](#) on page 339
- [“Setting the internal name of the clustered instance”](#) on page 343

Installing SQL Server 2005 on the second node

Before installing Microsoft SQL Server 2005, complete the preparation steps. See [“Preparing to install SQL Server 2005 on the second node”](#) on page 336.

Then complete the procedure to install and configure SQL Server 2005 on the second node.

SQL Server 2005 must have the same configuration on all nodes in the cluster. You will need the following information:

- Instance name (if applicable)
- Destination Folder for Program Files and Data Files
- Authentication Mode

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Multiple instances of SQL Server 2005 must be installed in the same order on every node of the cluster.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

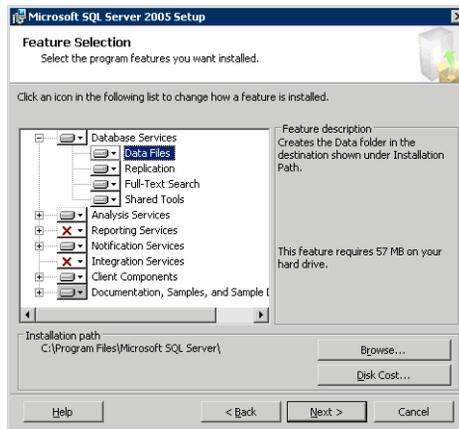
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

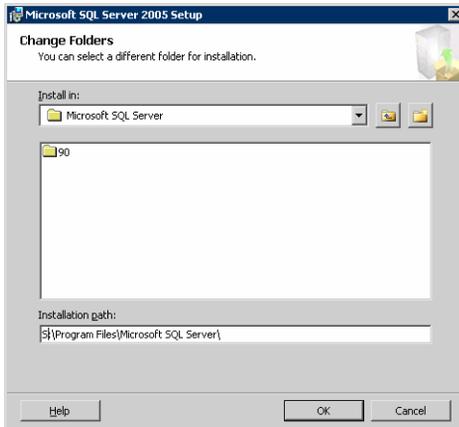
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.



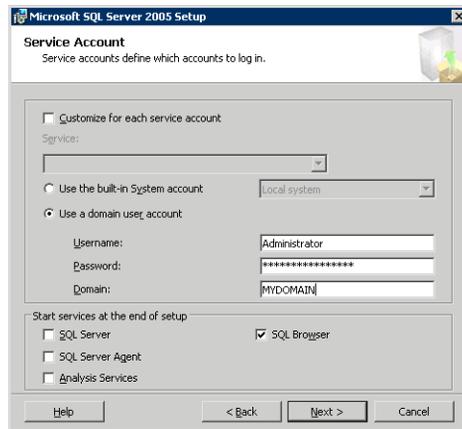
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 340, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.

Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.

- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.

See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, `INST1-VS\INST1`).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.

- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

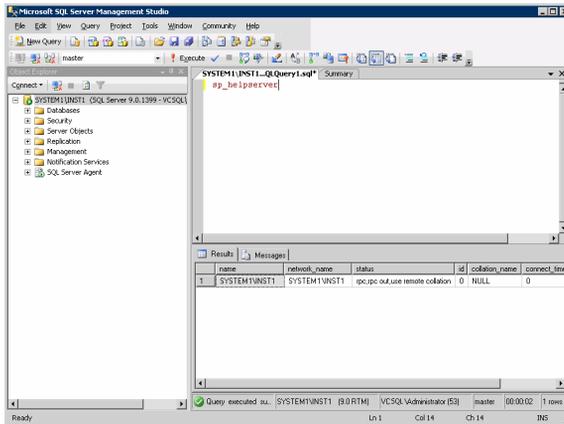
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, *SYSTEM1 \ INST1*. (For a default instance, you see only *System_Name* .)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
`sp_dropserver "System_Name\Instance_Name"`
 where **System_Name\Instance_Name** is the name noted in [step 3](#) on page 344.
 For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
 For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press **F5**.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.

Configuring the VCS SQL Server 2005 service group

A VCS SQL Server service group is used to bring a SQL Server 2005 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Prerequisites for configuring the service group

Complete the following tasks before configuring the service group:

- Verify that SFW HA, along with the VCS database agent for SQL Server 2005, is installed on all cluster nodes. See ["Installing Veritas Storage Foundation for Windows"](#) on page 304.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See ["Configuring the cluster"](#) on page 320.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- To avoid having to modify the service group later, create any new user-defined databases before creating the service group. See ["Adding a new SQL Server user-defined database"](#) on page 353.

- Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group. See “[Importing the cluster disk group](#)” on page 337 and “[Adding drive letters to mount the volumes](#)” on page 338.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. You specify the virtual server name when configuring the service group. See “[Setting the internal name of the clustered instance](#)” on page 343.

Note: For a disaster recovery configuration, the SQL Server virtual server name on the secondary site cluster must match the one on the primary site cluster.

- Assign a unique virtual IP address for the SQL Server 2005 instance. You specify this IP address when configuring the service group.
- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes.
 A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2005\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2005 Server service for the SQL instance.

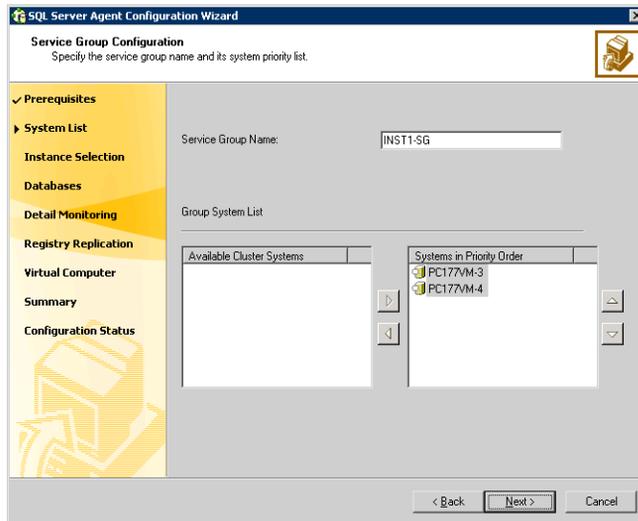
Creating the SQL Server 2005 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

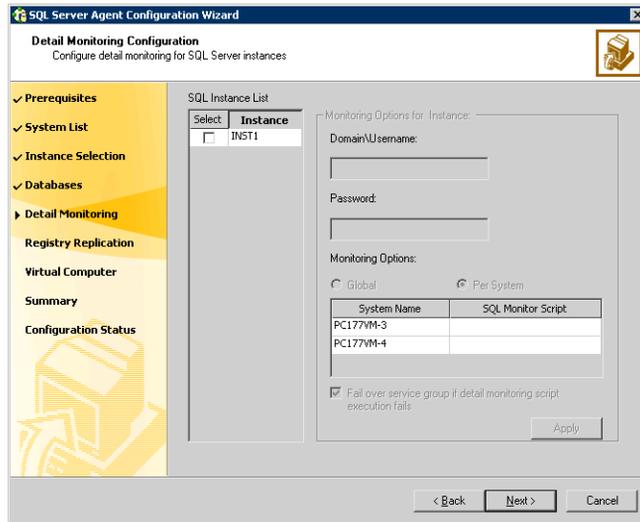
- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.

5 Specify the service group name and system list:



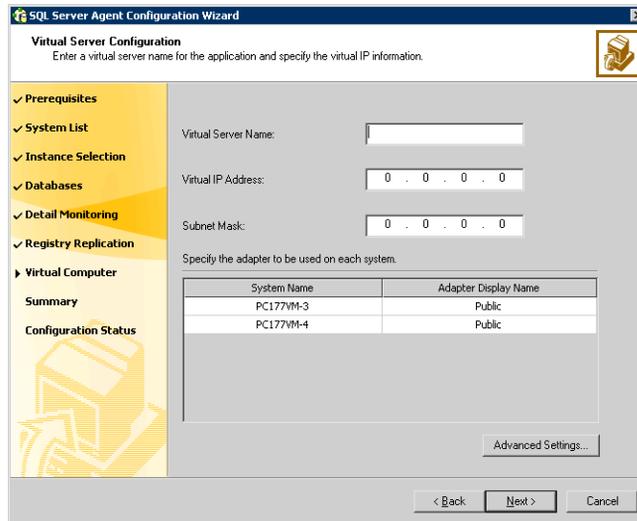
- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 348 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 349 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.

- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for high availability”](#) on page 355.

Verifying the SQL Server 2005 cluster configuration

Verify your installation by switching online nodes or by shutting down the computer that is currently online. Either process will test that the service group can be smoothly transferred between nodes.

Shutting down a node creates an actual failure, stressing your system, but more truly testing its high availability than by switching nodes. If you do shut down the online computer in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another node.

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

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

To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in step 1.
- 3 To move all the resources back to the original node, repeat step 1 for each of the service groups.

To shut down an active cluster node

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

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

Adding a new SQL Server user-defined database

If you want to add a SQL Server user-defined database to an existing configuration, complete the procedure “[Creating a SQL Server user-defined database](#)” on page 271. The procedure includes the following tasks:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Add the VMDg and MountV resources for the user databases.

Additional instructions for disaster recovery

After completing the tasks for setting up a SFW HA environment for SQL Server 2005 on a primary site, you can optionally create a secondary or “failover” site for disaster recovery.

See [Chapter 17, “Deploying disaster recovery: New SQL Server 2005 installation”](#) on page 881.

Configuring an MSDTC service group for high availability

This chapter includes the following topics:

- [“Tasks for configuring MSDTC for high availability”](#) on page 355
- [“Reviewing the prerequisites”](#) on page 357
- [“Reviewing the configuration”](#) on page 357
- [“Configuring cluster disk groups and volumes”](#) on page 360
- [“Mounting volumes used by the MSDTC service group”](#) on page 366
- [“Creating an MSDTC service group”](#) on page 367
- [“Creating an MSDTC client”](#) on page 369

Tasks for configuring MSDTC for high availability

The VCS Database Agent for SQL, which you install during SFW HA installation for Microsoft SQL high availability, includes an MSDTC agent. The MSDTC agent can provide high availability for the Microsoft Data Transaction Coordinator (MSDTC) service. The MSDTC agent comprises two parts: MSDTC client and MSDTC server.

You can configure high availability for MSDTC either before or after configuring high availability for Microsoft SQL. To configure high availability for MSDTC, you use the SQL Server Configuration Wizard to configure an MSDTC service group for the MSDTC server and to configure the MSDTC client.

To modify an existing MSDTC service group, see the *Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide*.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 8-1 Tasks for configuring MSDTC for high availability

Objective	Tasks
“Reviewing the prerequisites” on page 357	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 357	<ul style="list-style-type: none"> ■ Understanding MSDTC service group configuration ■ Reviewing the sample configuration
“Configuring cluster disk groups and volumes” on page 360	<ul style="list-style-type: none"> ■ Configuring cluster disk groups for an MSDTC service group ■ Configuring volumes for an MSDTC service group
“Mounting volumes used by the MSDTC service group” on page 366	Mounting the drives used by the service group on the node on which you are configuring the service group
“Creating an MSDTC service group” on page 367	Creating an MSDTC service group
“Creating an MSDTC client” on page 369	Creating an MSDTC client.

Reviewing the prerequisites

You must meet the following prerequisites before creating and configuring the MSDTC service group:

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that the VCS Database Agent for SQL is installed on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Cluster Configuration Wizard (VCW).
- Verify that the drives containing the MSDTC logs and registry replication directory are mounted on one node (the node on which you are configuring the service group) and unmounted on all other nodes.
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- Assign a unique virtual server name and virtual IP address for the MSDTC server.
- Verify that the Distributed Transaction Coordinator service is stopped.

Reviewing the configuration

MSDTC servers can coexist with SQL servers on the same cluster nodes. If the MSDTC server and the SQL server are running on the same node, the MSDTC client is left in the default configuration. If the MSDTC server is not configured on the same node as the SQL server, then the MSDTC client must be configured on that node using the SQL Server Configuration Wizard. The MSDTC client and the MSDTC server must not run on the same cluster node.

For instance, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage. The shared storage can be managed using Veritas Storage Foundation for Windows (SFW). Two configurations are possible:

- SQL server is configured on different nodes than the MSDTC server
- SQL server is configured on the same node as the MSDTC server

Figure 8-1 MSDTC server configured on different nodes than SQL server

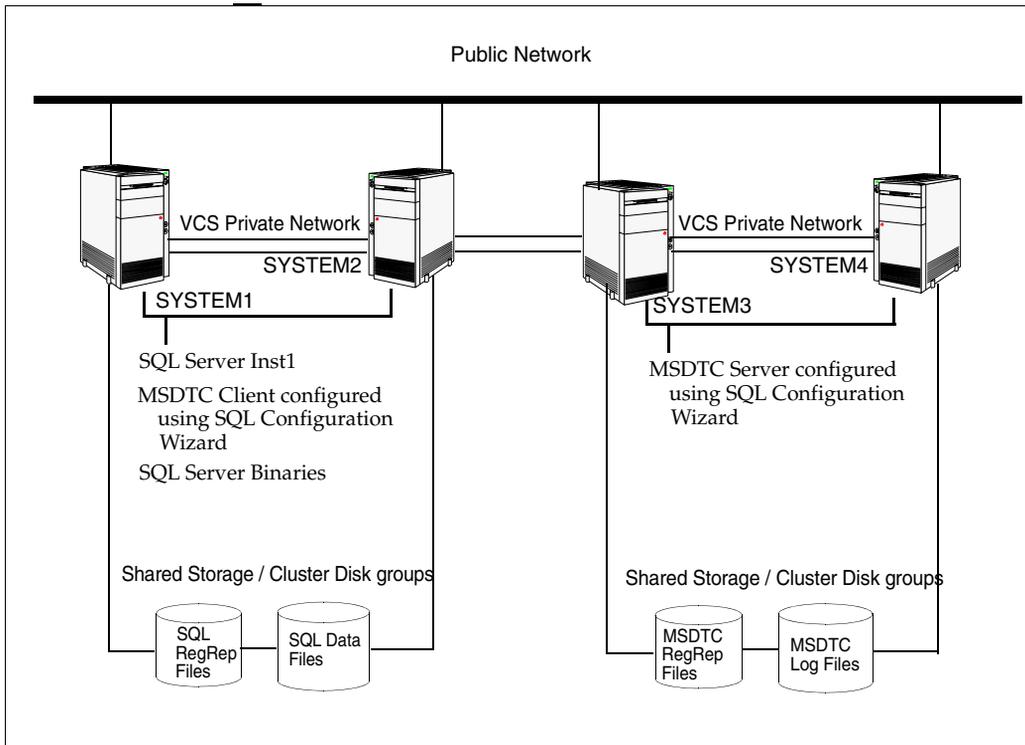
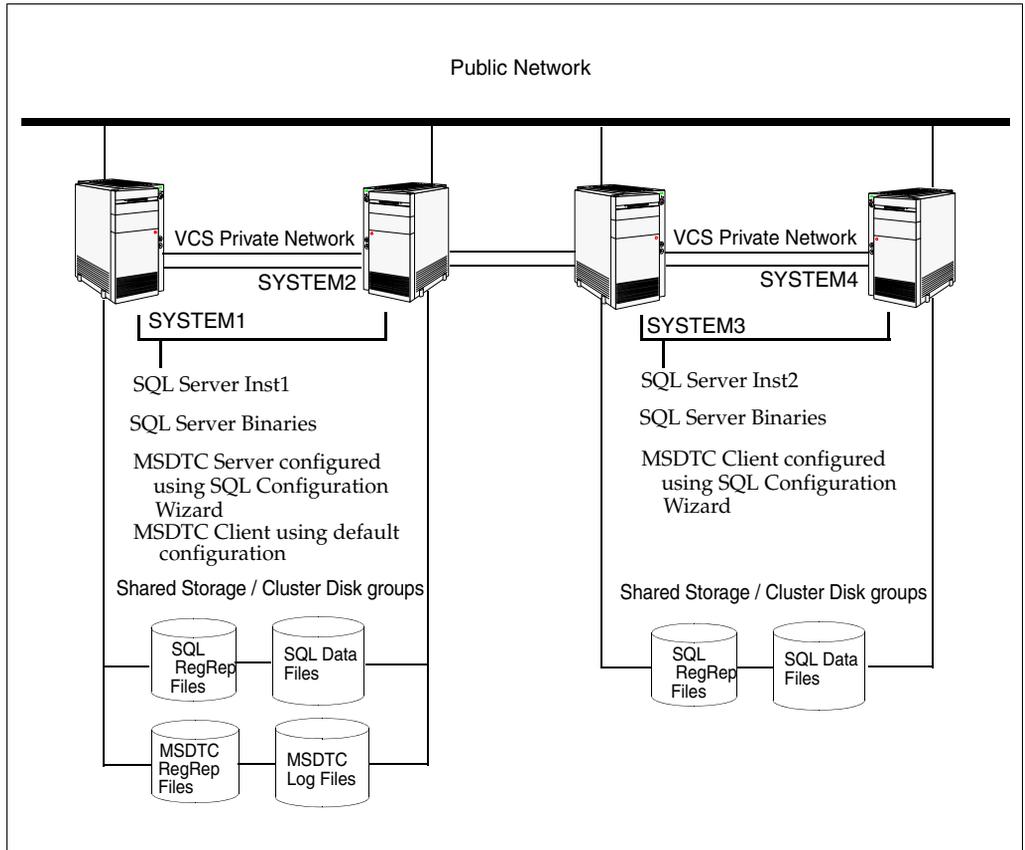


Figure 8-2 MSDTC server configured on the same node as SQL server



Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your MSDTC service group.

A disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Complete the following tasks before you create the cluster disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place MSDTC files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (MSDTC_DG) on shared disks and then create the following volumes:

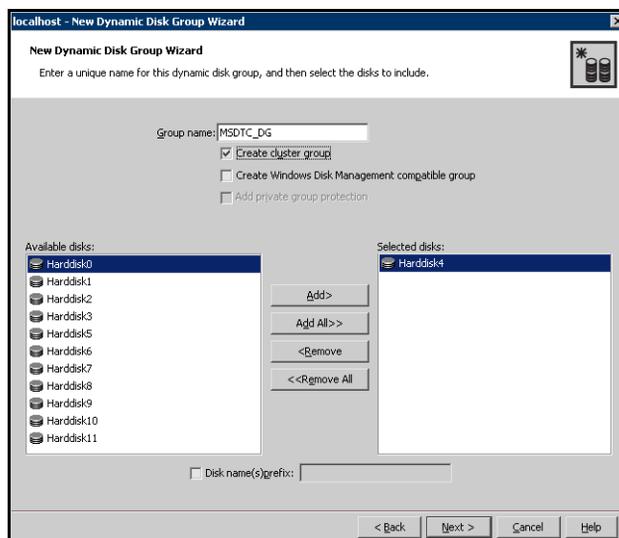
- MSDTC_LOG: contains the MSDTC log files.
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group. Create a 100 MB volume for this purpose.

Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

To create a cluster disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 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 On the **Welcome** page of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group you want to create.



- Enter the name of the disk group (for example, MSDTC_DG).
 - Click **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and click the **Add** button to move them to the **Selected disks** list.
 - Click **Next**.
- 7 Review the selected disks and click **Next**.
 - 8 Review the summary information and click **Finish**.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

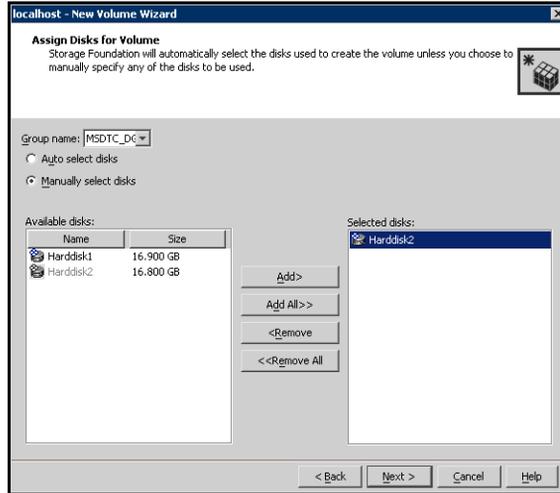
- MSDTC_LOG: contains the MSDTC log files
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group.

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

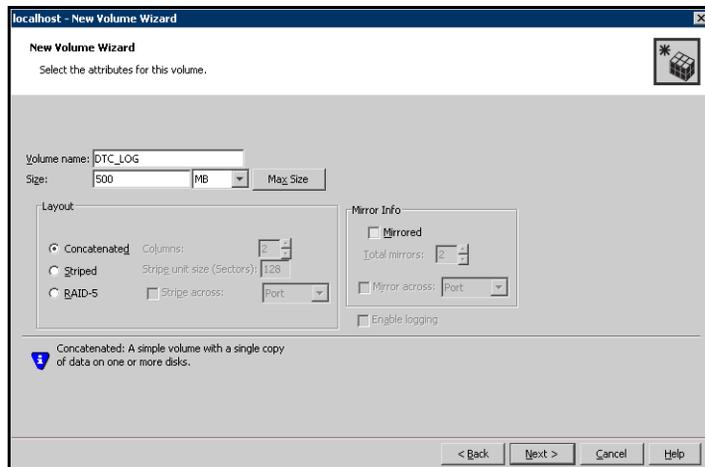
To create a volume

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the 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 the disk group you created for the MSDTC volumes and select **New Volume**.
- 5 On the **Welcome** page of the New Volume Wizard, click **Next**.

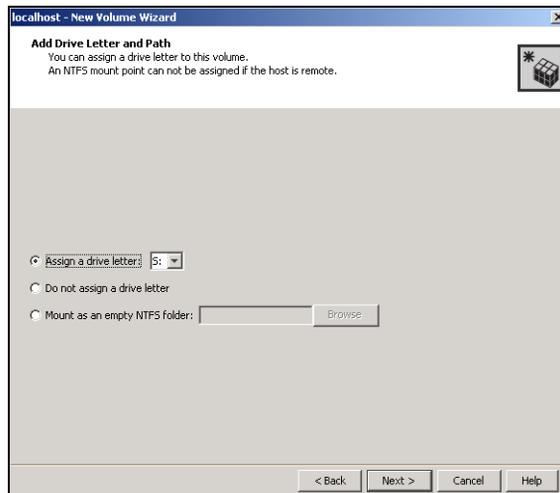
- 6 On the **Assign Disks for Volume** page, make sure the name of the disk group that you created for MSDTC appears in the Group name field.



- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended. After selecting the necessary disks, click **Next**.
- 8 Specify the volume attributes:

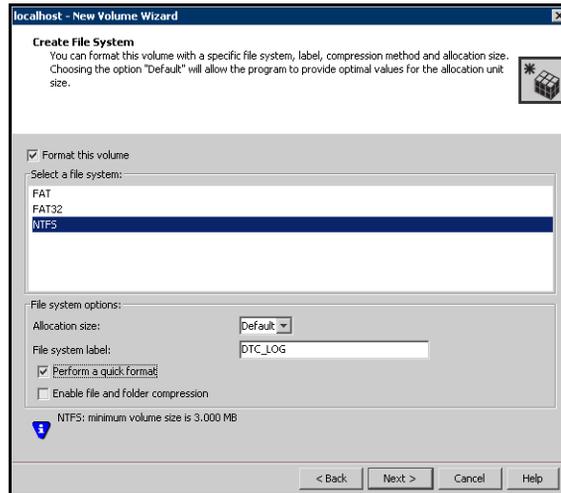


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, 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.
- 9 In the **Add Drive Letter and Path** page, 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.

- 10 Specify information about the attributes for the file system that will be created:



- 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**.
- 11 Review the Summary dialog box and click **Finish**.
 - 12 Repeat the procedure as necessary to create the additional volumes needed for your deployment. Create the volumes on the first node of the cluster only.

Mounting volumes used by the MSDTC service group

Ensure that the shared volumes created for the MSDTC log files and registry replication information are mounted on the node where you run the SQL Server Configuration Wizard to configure the service group. The volumes must be unmounted on all other nodes in the cluster.

If you created the MSDTC disk group and volumes on another node, mounting a volume involves deporting the disk group from the other node and then importing the disk group to the node where you are configuring the service group.

Occasionally, when a disk group is imported, a drive letter may not be associated with an existing volume. If this occurs, use VEA to add a drive letter or folder path and mount the volume so that it can be seen by the operating system.

To mount volumes

- 1 If the disk group is not imported, deport it from the other node and then import it on the current node:
 - Stop all processes accessing the volumes in the disk group, and from the VEA console's tree view, right-click the disk group and select **Deport Dynamic Disk Group**. Then select **Yes**.
 - In the VEA, connect to the node where you want to import the cluster disk group. In the tree, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node. Then right-click the disk group and select **Import Dynamic Group** and click **OK**.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths dialog box, select **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder. Assign the same drive letter or mount path that was assigned when the volume was created.
 - *To assign a drive letter:* In the Assign Drive Letter dialog box, select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder:* Select the **Mount as an empty NTFS folder option** and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**. Repeat [step 2](#) on page 366 through [step 5](#) on page 366 for all the volumes to be mounted. You can now create the MSDTC service group.

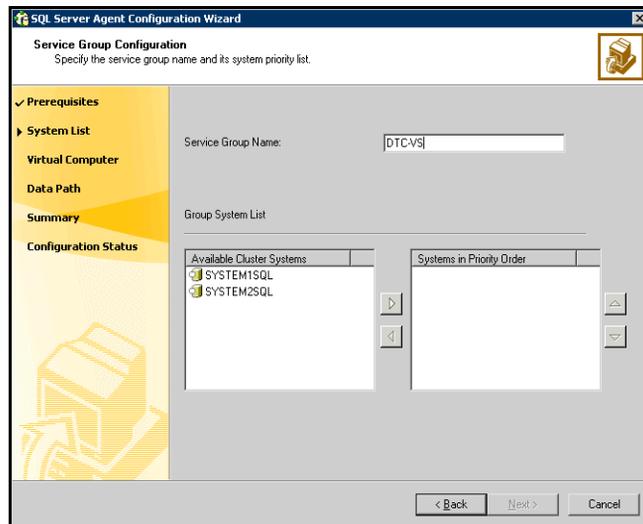
Creating an MSDTC service group

MSDTC is a global resource and can be accessed by more than one SQL Server service group. Symantec recommends that you configure one MSDTC service group in a VCS cluster.

Note: You must be a Local Administrator on the node where you run the wizard.

To configure an MSDTC service group

- 1 If you have just configured a SQL service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 2 Select **MSDTC Server - Service Group Configuration** and **Create**. Click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the service group name and system list.

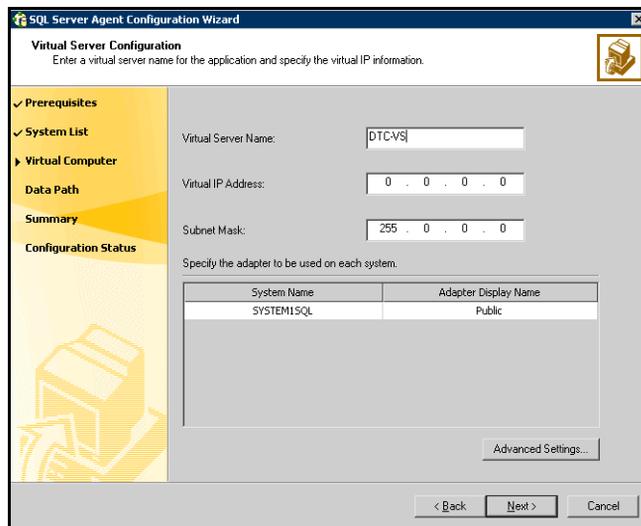


- Enter a name for MSDTC service group.
- In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right-arrow to move the systems to the service group's system list. Make sure you select the

systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.

- To change a system's priority in the **Systems in Priority Order** list, select the system and click the up and down arrows. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
- Click **Next**. If the configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.

5 Specify the information related to the virtual server.



- Enter a virtual server name for the node on which the DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.
- Enter a unique virtual IP address for the MSDTC server.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system.

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

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop down list.
 - Click **Next**.
- 6 Specify the MSDTC log and replication directory and click **Next**. Symantec recommends using different paths for these directories. If the directory does not exist, the wizard creates it.
 - 7 Review the service group configuration, change the resource names if desired and click **Next** to create the service group.
 - The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.
 - The wizard assigns unique names to resources. Change names of the resources, if desired. To edit a resource name, select the resource name and either click it or press the **F2** key. Press **Enter** after editing each resource name. To cancel editing a resource name, press **Esc**.
 - 8 A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes** to create the service group.
Various messages indicate the status of these commands.
 - 9 In the Configuration Complete panel, check **Bring the service group online** to bring the configured service group online. To bring the service group online later, uncheck the option.
 - 10 Click **Next** to create an MSDTC client, or click **Finish** to exit the wizard.

Creating an MSDTC client

Set the MSDTC client to run on nodes where:

- A SQL instance is configured to run
- The MSDTC server is not configured to run.

Before running the wizard to configure an MSDTC client:

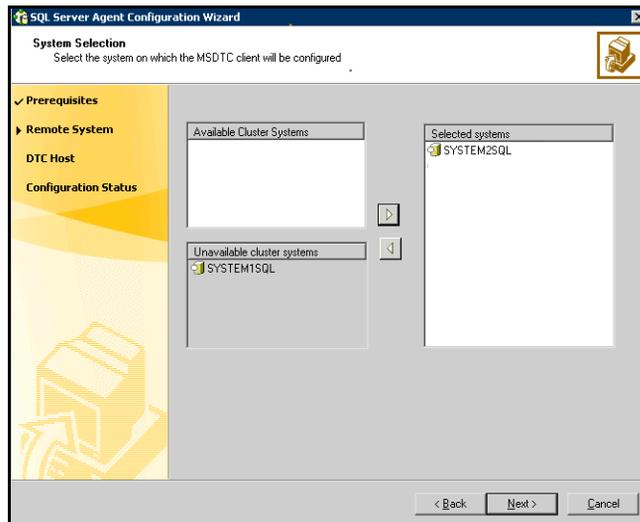
- Verify that the MSDTC service group is online in the cluster.
- Verify that the node on which you run the wizard is not a part of an MSDTC service group system list.

To configure an MSDTC client

- 1 If you have just configured the MSDTC service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL

Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.

- 2 In the Select Configuration Option panel, select **MSDTC Client Configuration** and click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 In the System Selection panel, specify the nodes on which the MSDTC client will be configured and click **Next**:



- Select the nodes in the **Available Cluster Systems** list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - Click the right arrow to add them to the **Selected systems** list.
 - The **Unavailable Cluster Systems** lists the nodes that have an MSDTC service group configured and are therefore not available for setting up an MSDTC client.
- 5 If the MSDTC service group is not online in the cluster, an informational message appears informing you that the wizard will bring the MSDTC service group online. Click **Yes**.
 - 6 On the Specify DTC Node panel, specify the virtual DTC server name and click **Next**.
 - 7 On the Configuration Complete panel, click **Finish** to exit the wizard.

Campus Cluster

This section contains the following chapters:

- [Campus cluster for SQL Server: Overview](#)
- [Deploying SFW HA for Campus Cluster: New SQL Server 2000 Installation](#)
- [Deploying SFW HA for Campus Cluster: New SQL Server 2005 Installation](#)

Campus cluster for SQL Server: Overview

This chapter includes the following topics:

- [“What is a campus cluster?”](#) on page 374
- [“Differences between campus clusters and local clusters”](#) on page 374
- [“Why implement a campus cluster?”](#) on page 376
- [“Campus cluster failover using the ForceImport attribute”](#) on page 376

What is a campus cluster?

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 with the same number of disks and contains mirrored data of the storage on the other array.

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.

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.

Administrators can use campus clusters to protect data from natural disasters, such as floods and hurricanes, and unpredictable power outages. Campus clusters provide a layer of protection that extends beyond local high availability but is not as complex as disaster recovery with replication.

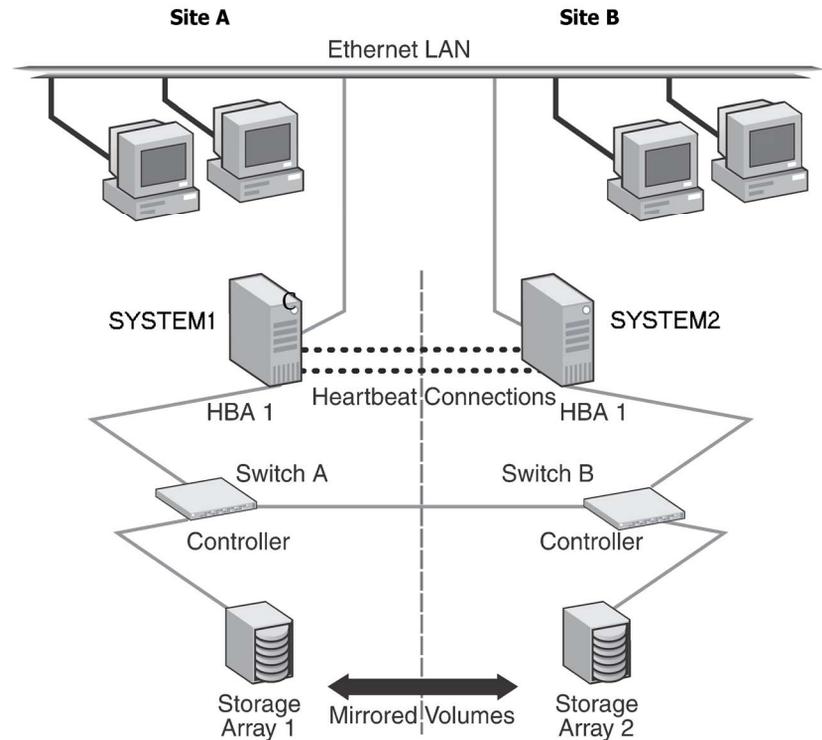
Differences between campus clusters and local clusters

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.

Sample Campus Cluster configuration

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

Figure 9-1 SQL Campus Cluster: Active-Passive configuration



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

If a site failure occurs in a two-node campus cluster, the remaining cluster node will not be able to bring the cluster disk groups online because it cannot reserve a majority of disks in the disk groups. To allow for failover to the other site, a

procedure forces the import to the other node, allowing a cluster disk group to be brought online on another node when that node has a minority of the cluster disks.

Implementing these force import procedures should be done with care. The primary site may appear to have failed but what really has happened is that both the storage interconnect between sites and the heartbeats have been lost. In that case, cluster disk groups can still be online on the primary node. If a force import is done so that the data can be accessed on the secondary site, the cluster disks will be online on both sites, risking data corruption.

Why implement a campus cluster?

In the event of a site disaster, such as power failure in a building, campus clusters offer a level of high availability that surpasses mirroring or clustering at a single site by dispersing the clustered servers into different buildings or sites. This environment also provides a simpler solution for disaster recovery than a more elaborate SFW HA DR environment with replication software; however, a campus cluster generally stretches a shorter distance than a replication-based solution depending on the hardware.

Campus cluster failover using the ForceImport attribute

Automated recovery is handled differently in a VCS campus cluster than with a VCS local cluster. [Table 9-1](#) lists failure situations and the outcomes depending on the settings for the ForceImport attribute of the VMDg resource. You can set this attribute to 1 (forcing the import of the disk groups to the other node) or 0 (not forcing the import). Use the VCS Java Console or command line to modify the ForceImport attribute.

To ensure proper failover in a VCS campus cluster, you must verify the value of the ForceImport attribute of the VMDg resource.

Table 9-1 Failure Situations

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic force import)
<p>1) Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.</p>	<p>Application automatically moves to another node.</p>	<p>Service Group failover is automatic on the standby or preferred system or node.</p>
<p>2) Server failure May mean a power cord became unplugged or a failure caused the system to stop responding.</p>	<p>Application automatically moves to other node. 100% of the disks are still available.</p>	<p>Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.</p>
<p>3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.</p>	<p>No interruption of service. Remaining disks in mirror are still accessible from the other node.</p>	<p>The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.</p>
<p>4) Zone failure Complete Site failure, all accessibility to the servers and storage is lost.</p>	<p>Manual intervention required to online the Service Group at remaining site. Can not automatically import 50% of mirrored disk.</p>	<p>Automatic failover of Service Group to online site. Force Import must be set to True before site failure to ensure VCS can import 50% of mirrored disk.</p>
<p>5) Split-brain (loss of both heartbeats) If the public network link serves as a low-priority heartbeat, the assumption is made that the link is also lost.</p>	<p>No interruption of service. Can't import disks because the original node still has the SCSI reservation.</p>	<p>No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.</p>

Table 9-1 Failure Situations (Continued)

Failure Situation	Forcelmport set to 0 (import not forced)	Forcelmport set to 1 (automatic force import)
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 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. Cannot import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.	Automatically imports 50% of mirrored disk to the alternate node. Disks online for a short period in both locations but offlined again due to IP and other resources being online on original node. No interruption of service.

Reinstating faulted hardware

Once a failure occurs and an application is migrated to another node or site, it is important to know what will happen when the original hardware is reinstated. For failure scenarios 3 through 7 described in [Table 9-1](#) on page 377 earlier, [Table 9-2](#) on page 378 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 [Table 9-1](#) on page 377 earlier.

Table 9-2 Behavior exhibited when hardware is reinstated

Failure Situation, before Reinstating the Configuration	Forcelmport set to 0 (import not forced)	Forcelmport set to 1 (automatic force import)
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Resync the mirror from the remote site.	Same behavior.

Table 9-2 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)
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.

Deploying SFW HA for Campus Cluster: New SQL Server 2000 Installation

This chapter covers the following topics:

- [“Reviewing the requirements”](#) on page 385
- [“Reviewing the configuration”](#) on page 390
- [“Configuring the storage hardware and network”](#) on page 391
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 394
- [“Configuring the cluster”](#) on page 400
- [“Configuring cluster disk groups and volumes for SQL Server 2000”](#) on page 418
- [“Managing disk groups and volumes”](#) on page 427
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 429
- [“Preparing to install SQL Server 2000 on the second node”](#) on page 432
- [“Installing and configuring SQL Server 2000 on the second node”](#) on page 436
- [“Setting the internal name of the clustered instance”](#) on page 440
- [“Creating a SQL Server user-defined database”](#) on page 442
- [“Completing configuration steps in SQL Server”](#) on page 443
- [“Configuring the SQL Server 2000 service group for VCS”](#) on page 444

- [“Modifying the IP resource in the SQL Server 2000 service group”](#) on page 450
- [“Verifying the campus cluster: Switching the service group”](#) on page 451
- [“Setting the ForceImport attribute to 1 after a site failure”](#) on page 451
- [“Modifying the SQL 2000 service group to add VMDg and MountV resources”](#) on page 452

This chapter provides information on how to install and configure a new Veritas Storage Foundation HA environment for SQL Server in a campus cluster. A campus cluster environment provides high availability and disaster recovery that extends beyond local clustering and mirroring at a single site, but is not as complex as SFW HA DR solution with replication.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2000.

See [“Using the Solutions Configuration Center”](#) on page 31.

Table 10-1 on page 383 outlines the high-level objectives and the tasks to complete each objective.

Table 10-1 Task list: SQL Server Campus Cluster configuration

Objective	Tasks
“ Reviewing the requirements ” on page 385	Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 390	<ul style="list-style-type: none"> ■ Understanding active-passive configuration ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 391	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Installing Veritas Storage Foundation HA for Windows ” on page 394	<ul style="list-style-type: none"> ■ Verifying the driver signing options for the system ■ Installing SFW and VCS (automatic installation) and installing Veritas Cluster Server Application Agent for Microsoft SQL Server ■ Restoring driver signing options for the system
“ Configuring cluster disk groups and volumes for SQL Server 2000 ” on page 418	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“ Configuring the cluster ” on page 400	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the Veritas Cluster Server Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“ Installing and configuring SQL Server 2000 on the first node ” on page 429	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Setting SQL Server services to manual start

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

Objective	Tasks
“Preparing to install SQL Server 2000 on the second node” on page 432	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Removing shared SQL files
“Installing and configuring SQL Server 2000 on the second node” on page 436	Installing and configuring SQL Server 2000
“Setting the internal name of the clustered instance” on page 440	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 442	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Configuring the SQL Server 2000 service group for VCS” on page 444	Creating a SQL Server service group using the SQL Server Configuration Wizard
“Modifying the IP resource in the SQL Server 2000 service group” on page 450	Modifying the Address and SubNetMask attributes if the sites are in different subnets.
“Setting the ForceImport attribute to 1 after a site failure” on page 451	If a site failure occurs, setting the ForceImport attribute of the VMDg resource to 1 to ensure proper failover.

Reviewing the requirements

The campus cluster solution allows for clustered systems with mirrored or synchronously replicated storage arrays to be implemented in separate datacenters, located either within the same building or separate buildings. For example, datacenter A could be located in building A and datacenter B located in building B. This guide will refer to these different areas as Site A and Site B.

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

Disk space requirements

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

[Table 10-2](#) on page 385 estimates disk space requirements for SFW HA.

Table 10-2 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
--	---

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 389.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.

- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

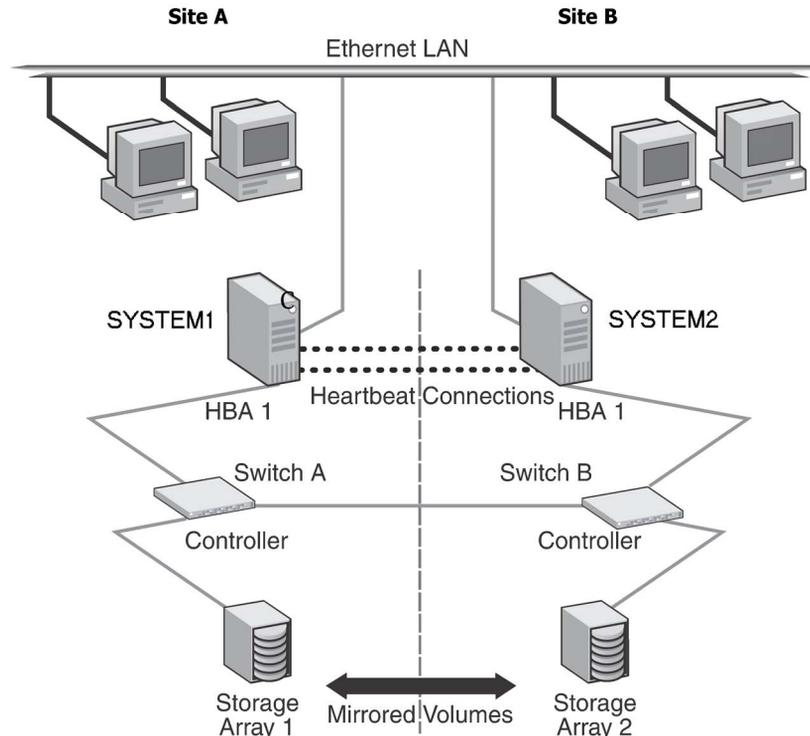
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

This chapter uses the example of a two-node campus cluster with each node in a separate site (Site A and Site B). In this example, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array.

Figure 10-1 illustrates a Active-Passive configuration for SQL Server with one to one failover capabilities. The active node of the cluster hosts the SQL 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. In this case, the SQL virtual server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

Figure 10-1 SQL Campus Cluster: Active-Passive configuration



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 should contain the same number of disks on each site for the mirrored volumes.

Following the workflow in the Solutions Configuration Center

The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA Campus Cluster environment for one or more instances of SQL Server 2000, in either an Active-Passive configuration.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.

- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 10-3](#) on page 394 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 10-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 395.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

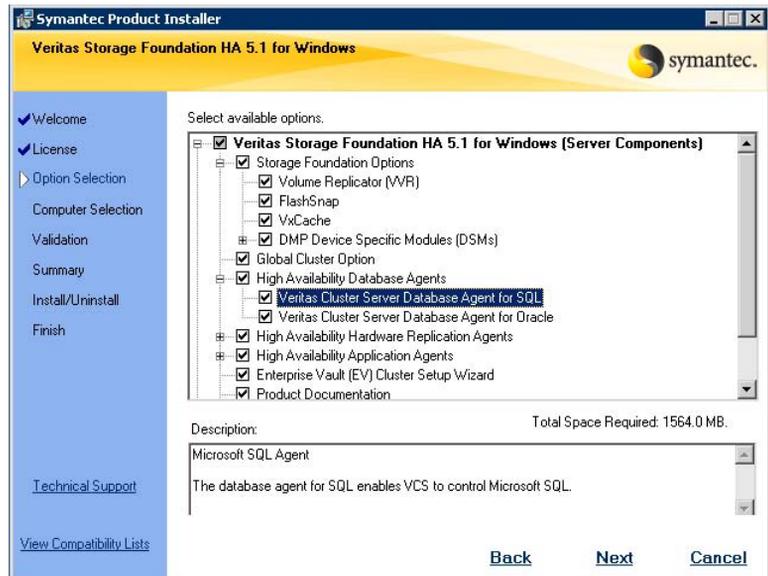
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL

Required to configure high availability for SQL Server.

Client

Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option

Required for a disaster recovery configuration only.

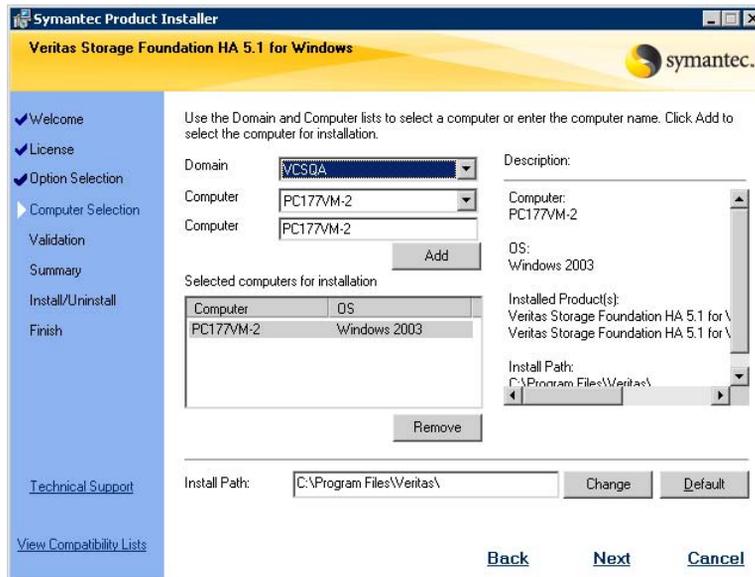
Veritas Volume Replicator

For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

High Availability Hardware Replication Agents

If you plan to use hardware replication, select the appropriate hardware replication agent.

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.

- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

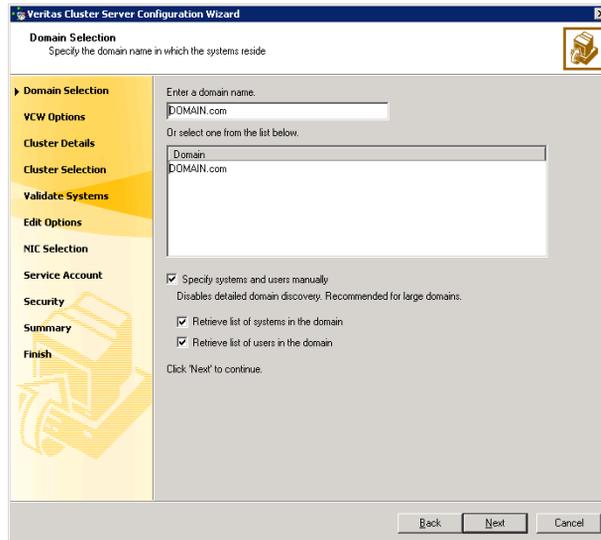
Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

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

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

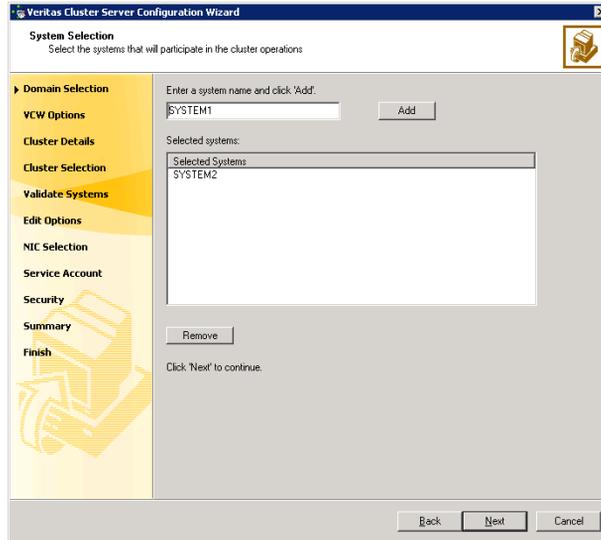
- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 8](#) on page 404.

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

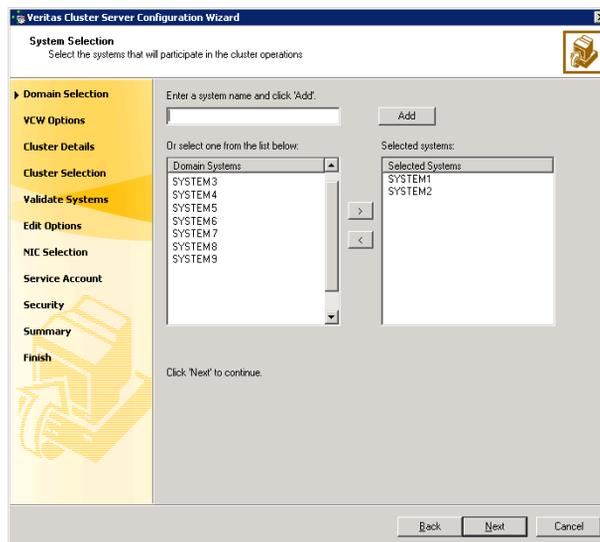
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 403. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 404.

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



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

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

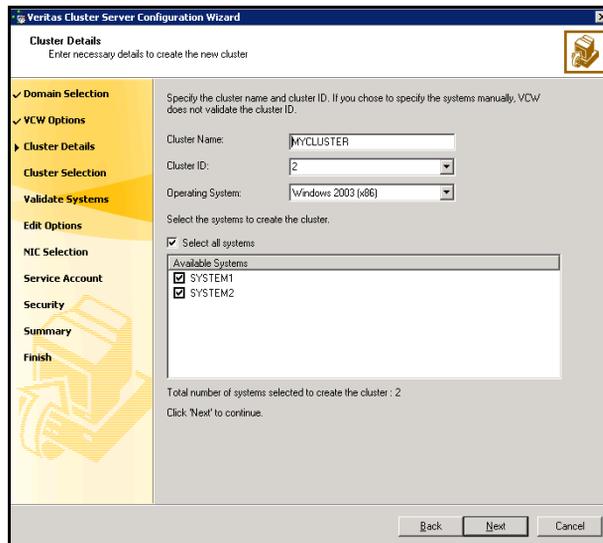
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

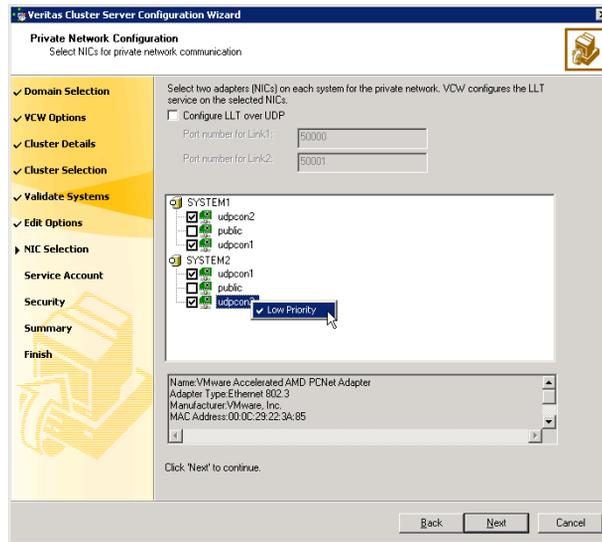
- On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- On the Cluster Details panel, specify the details for the cluster and then click **Next**.



Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255. Caution: If you chose to specify systems and users manually in step 4 on page 402 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 that the systems are running.
Available Systems	Select the systems that will be part of the cluster. 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. Check the Select all systems check box to select all the systems simultaneously.

- 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](#) on page 404, proceed to the next step. Otherwise, proceed to [step 12](#) on page 408.
- 11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.
 Do one of the following:

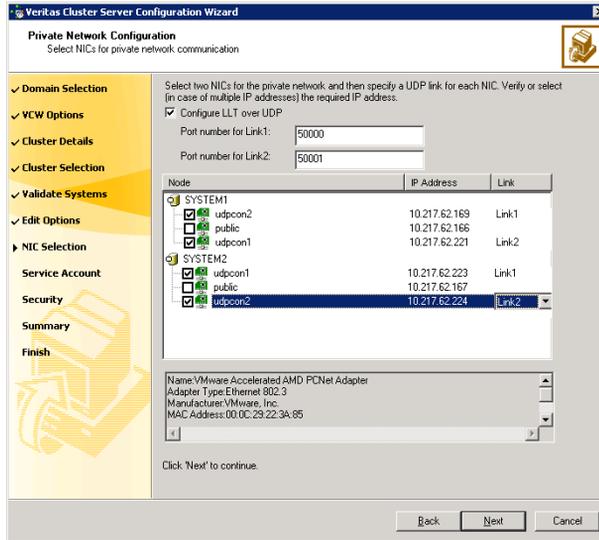
- To configure the VCS private network over Ethernet



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

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

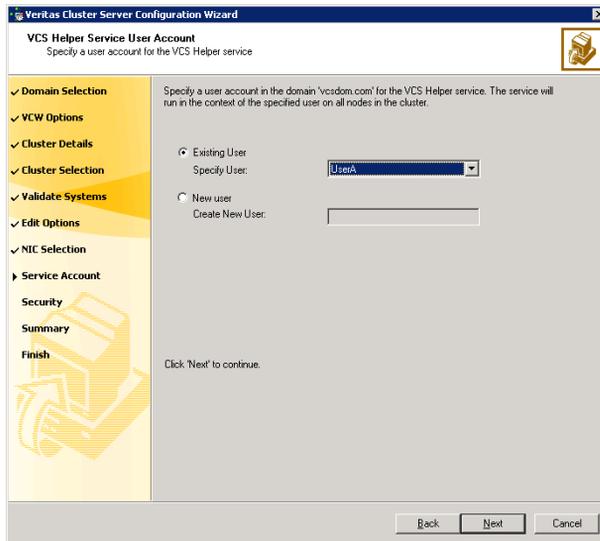
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



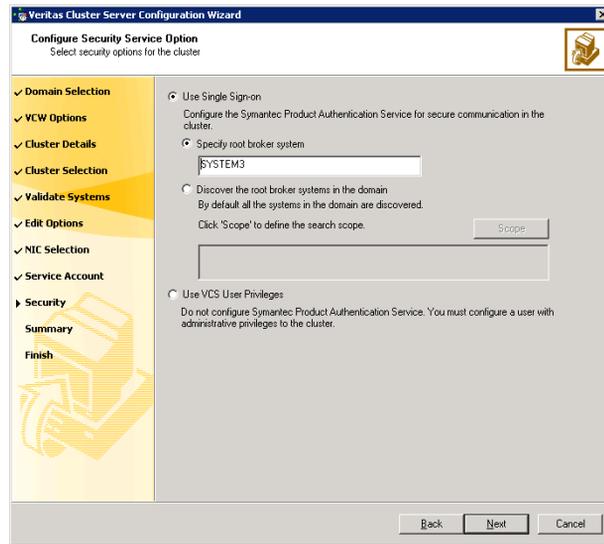
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 402, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 10-4](#) on page 410 contains some more examples of search criteria.

Table 10-4 Search criteria examples

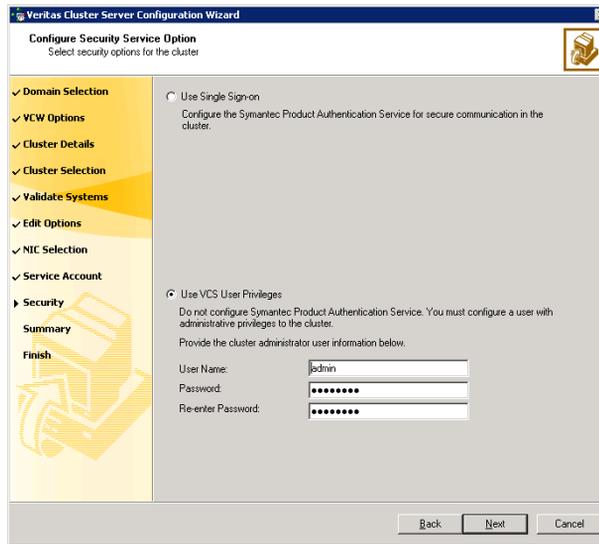
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for

the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

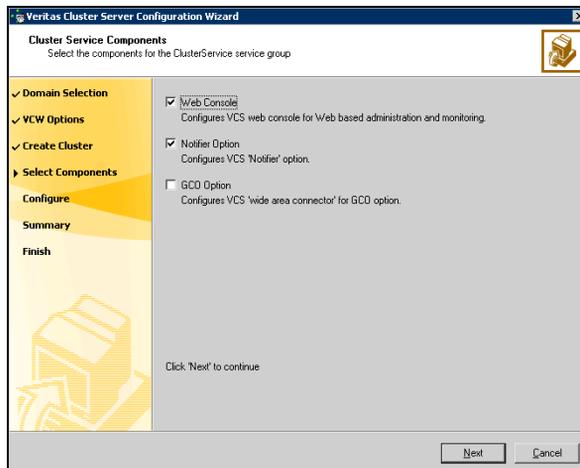
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



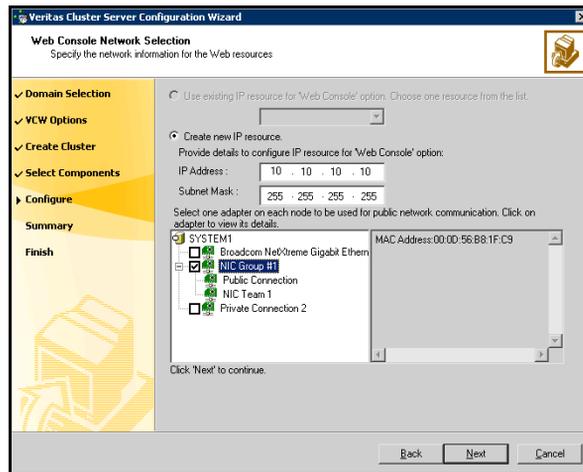
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 413.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 414.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

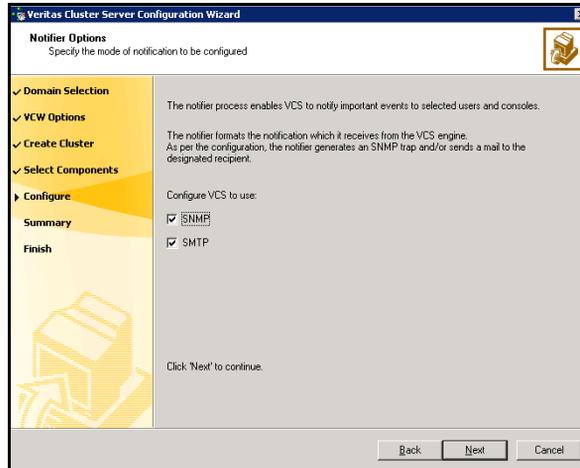
- 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 414. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

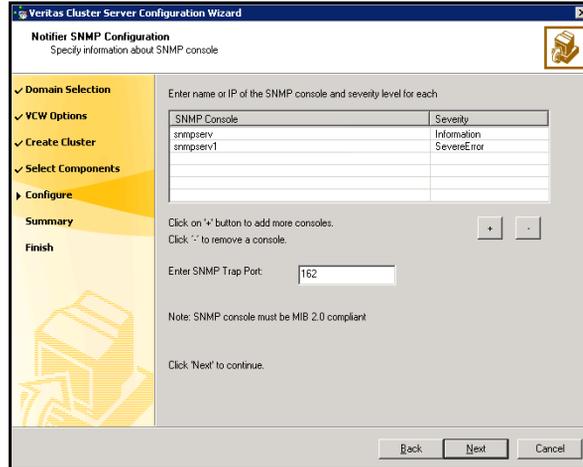
To configure notification

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



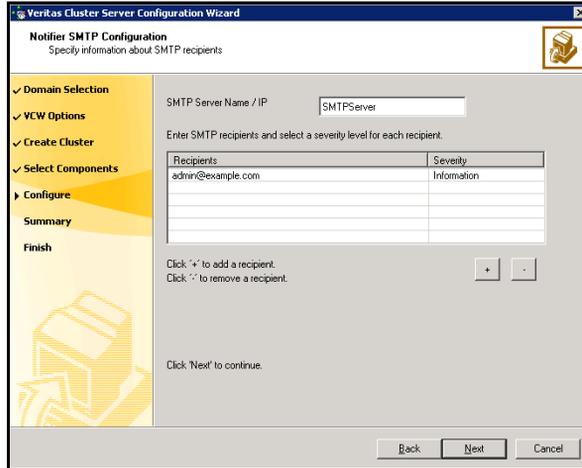
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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



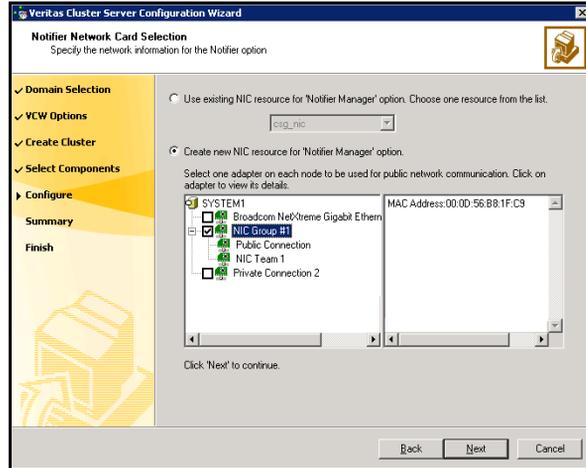
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



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

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



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

Configuring cluster disk groups and volumes for SQL Server 2000

Before installing SQL Server, you must create disk groups and mirrored volumes using the VEA console installed with SFW. This is also an opportunity to increase the size of existing volumes, add storage groups, and create volumes to support additional databases for storage groups.

About cluster disk groups and volumes

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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 fail over 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: Create a cluster disk group and volumes on only one node of a cluster. You make the volumes accessible to 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the *Veritas Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring cluster disk groups and volumes

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

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.

Complete the following tasks before you create the cluster disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size. Also, consider
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node of the cluster where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

Configuring the disks and volumes

Ensure that each disk group has the same number of disks on each site. Each volume must be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 421
- [“Creating a volume”](#) on page 423

Considerations when creating new volumes

Consider the following when creating new volumes.

- For campus clusters, when creating a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored 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.
- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.

- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
 To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
 The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

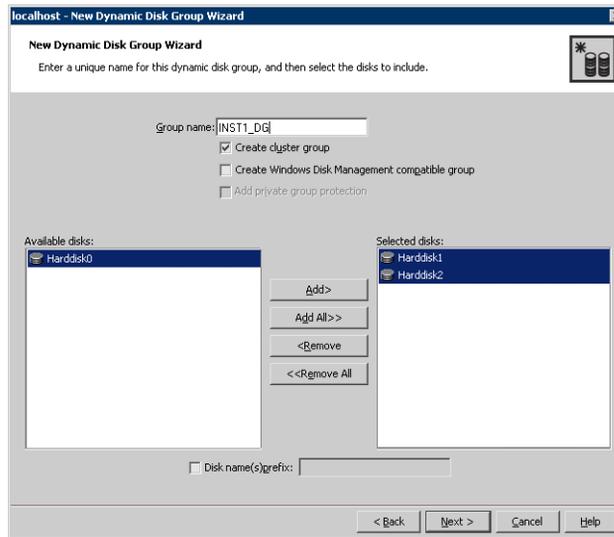
Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Proceed to create the appropriate volumes on each disk.

Creating a volume

Use this procedure to create the following volumes on the first node of the cluster:

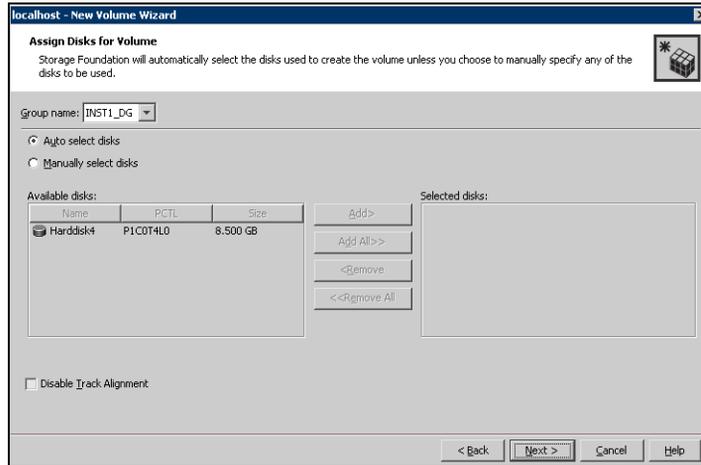
- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_DB1_VOL**: contains the user database files
- **INST1_DB1_LOG**: contains the user database log files
- **INST1_REGREP_VOL**: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

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

To create dynamic volumes

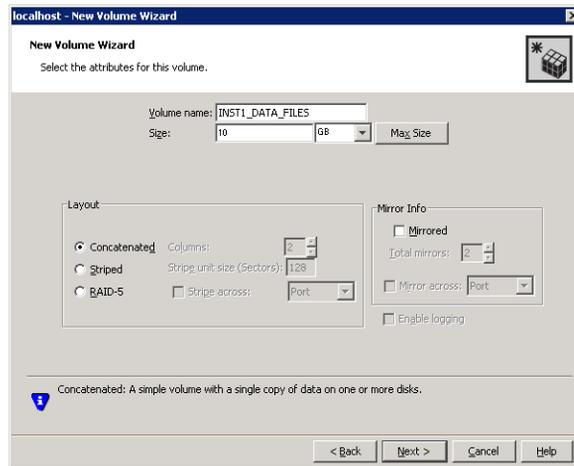
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example **INST1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



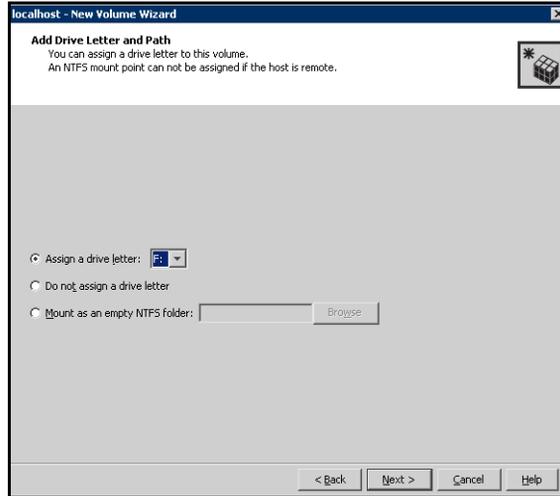
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the parameters of the volume.

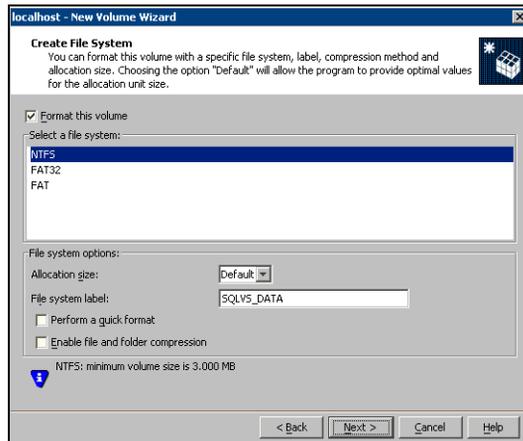


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation file size or accept the Default.

- The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.

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

Managing disk groups and volumes

This section includes the following procedures:

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

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

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Installing and configuring SQL Server 2000 on the first node

Complete the following tasks before installing SQL Server 2000:

- Verify that the cluster disk group is imported to the first node for this SQL instance
See “[Importing the cluster disk group](#)” on page 434.
- Mount the required volumes and ensure that they are assigned drive letter.
See “[Adding drive letters to mount the volumes](#)” on page 434.

Installing Microsoft SQL Server 2000

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.

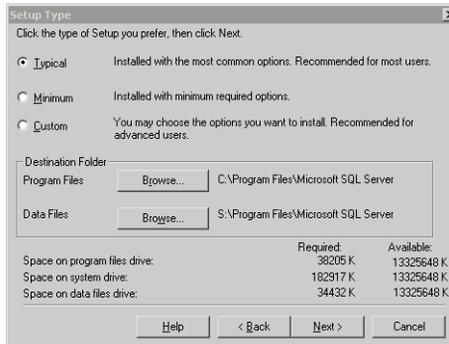
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

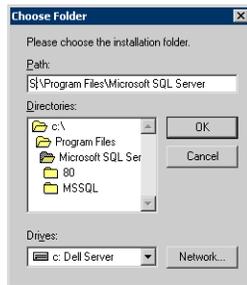
To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.

- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

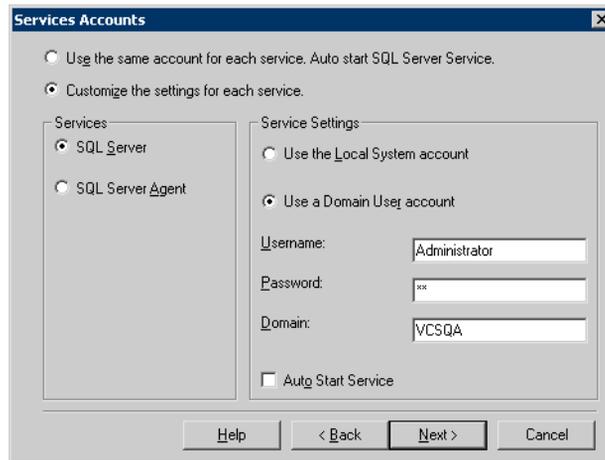


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

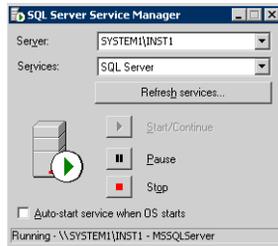
Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Preparing to install SQL Server 2000 on the second node

Complete the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:

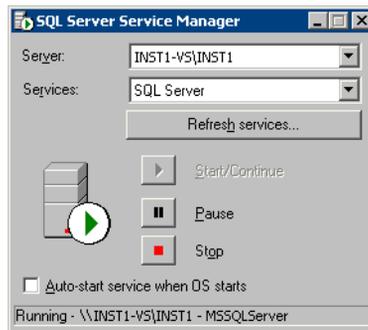
- [“Stopping the SQL Server 2000 service”](#) on page 433
- [“Deporting the cluster disk group”](#) on page 433
- [“Importing the cluster disk group”](#) on page 434
- [“Adding drive letters to mount the volumes”](#) on page 434
- [“Renaming shared SQL Server 2000 files”](#) on page 436

Stopping the SQL Server 2000 service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.

- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

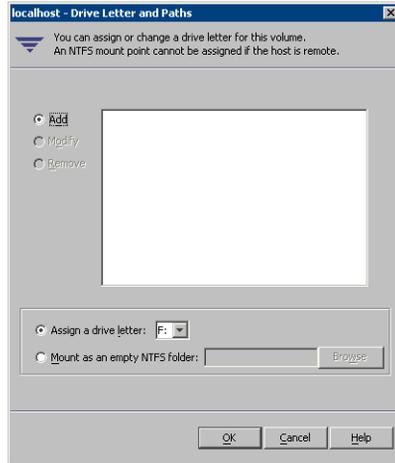
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2000 on the second node

Use the following procedures to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2000 on the second node”](#) on page 436
- [“Removing shared SQL Server 2000 files”](#) on page 439

Installing SQL Server 2000 on the second node

Before installing Microsoft SQL Server 2000:

- Verify that the cluster disk group is imported to the second node
See [“Importing the cluster disk group”](#) on page 434
- Verify that the volumes are mounted (are assigned drive letters)
See [“Adding drive letters to mount the volumes”](#) on page 434

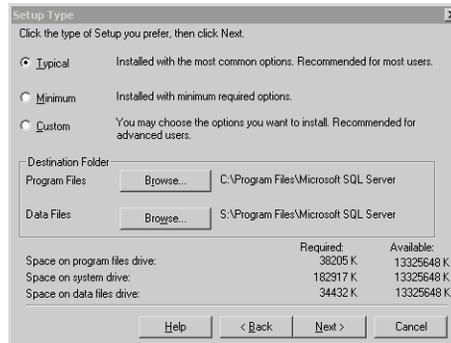
Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

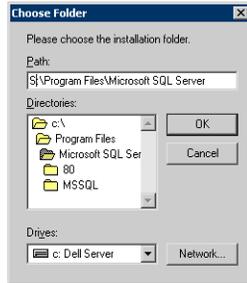
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

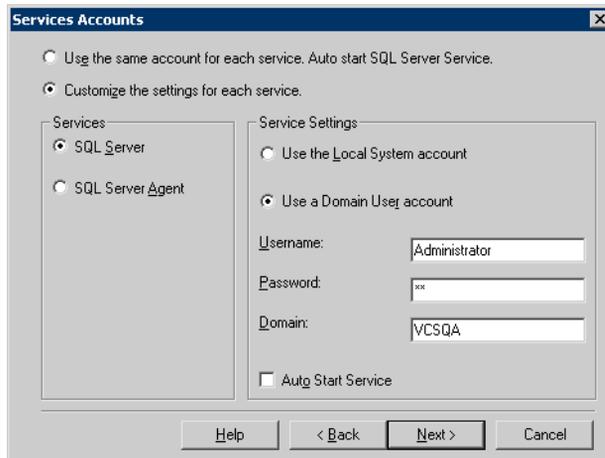
- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
- In the Services box, select the **SQL Server** option.
- In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
- Clear the **Auto Start Service** option.

- Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
 - 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

After you complete the SQL Server installation, repeat the following procedures on the additional nodes:

- [“Preparing to install SQL Server 2000 on the second node”](#) on page 432
- [“Installing and configuring SQL Server 2000 on the second node”](#) on page 436

Removing shared SQL Server 2000 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

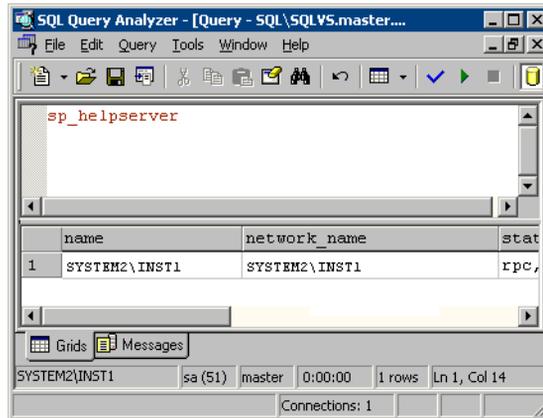
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`“sp_dropserver `System_Name\Instance_Name.’”`
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
`“sp_dropserver `SYSTEM2\INST1.’”`
 For example, for a default instance:
`“sp_dropserver `SYSTEM1.’”`
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
`“sp_addserver `Virtual_Server_Name\Instance_Name’,
 local”`

For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.

- Press F5.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database. If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See “[Creating a volume](#)” on page 423

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.
- 8 Configure whatever other options are required for your database.
- 9 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See “[Completing configuration steps in SQL Server](#)” on page 443.
- 10 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See “[Modifying the SQL 2000 service group to add VMDg and MountV resources](#)” on page 452.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 443.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating a volume](#)” on page 423.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Configuring the SQL Server 2000 service group for VCS

Configuring the SQL Server 2000 service group involves creating an SQL Server 2000 service group and defining the attribute values for its resources. A VCS SQL Server service group is used to bring a SQL Server 2000 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Prerequisites for configuring the SQL Server 2000 service group

- Verify that SFW HA, along with the VCS application agent for SQL Server 2000, is installed on all cluster nodes.
See “[Installing Veritas Storage Foundation HA for Windows](#)” on page 394.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW).
See “[Configuring the cluster](#)” on page 400.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group.
See “[Importing the cluster disk group](#)” on page 434.
See “[Adding drive letters to mount the volumes](#)” on page 434.
- Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name when configuring the service group.
See “[Setting the internal name of the clustered instance](#)” on page 440.
- Assign a unique virtual IP address for the SQL Server 2000 instance. You specify this IP address when configuring the service group.

- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists on all the cluster nodes. A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2000\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2000 Server service for the SQL instance.
See “[Stopping the SQL Server 2000 service](#)” on page 433.

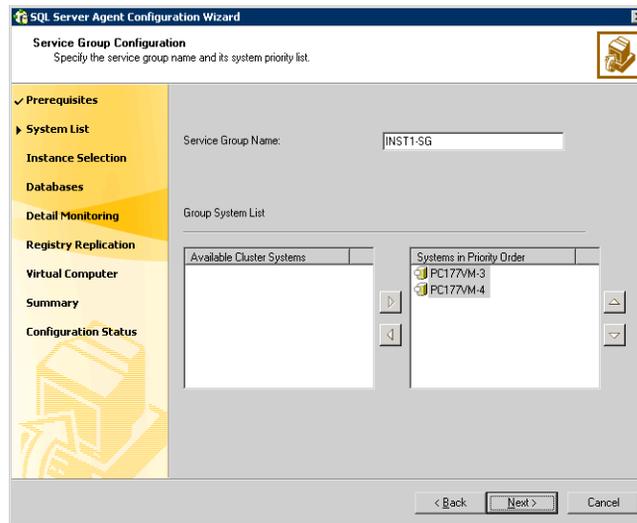
Creating the SQL Server 2000 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

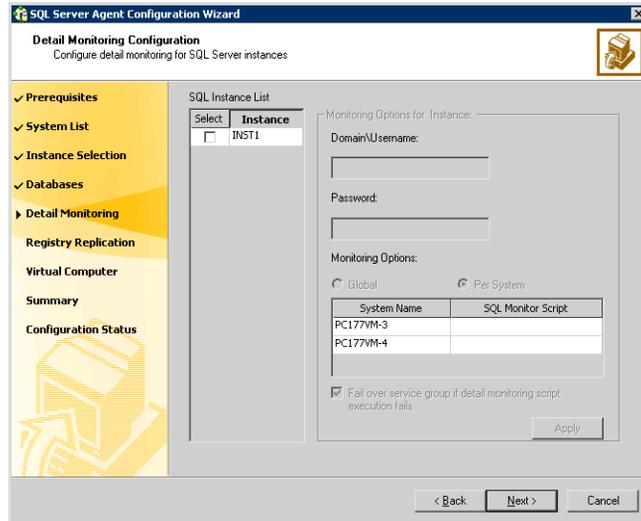
- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.

5 Specify the service group name and system list:



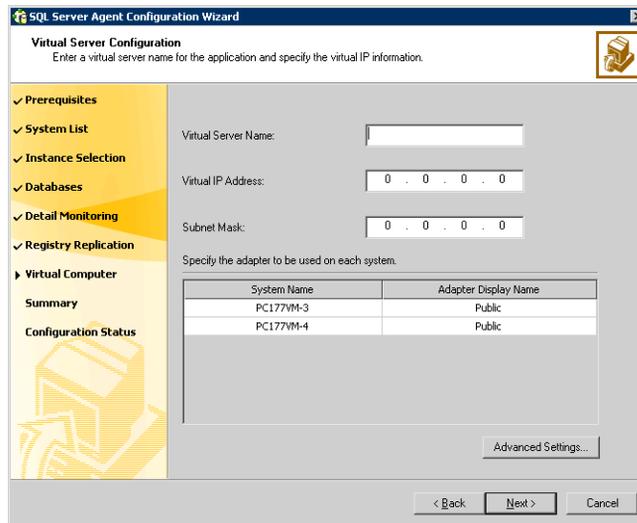
- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 446 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 447 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.

- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example `INST1-VS`. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.

- Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

If you have created a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

See “[Modifying the SQL 2000 service group to add VMDg and MountV resources](#)” on page 452.

Modifying the IP resource in the SQL Server 2000 service group

Note: This procedure is only applicable to a campus cluster with sites in different subnets.

Use the Java Console to modify the Address and SubNetMask attributes of the IP resource in the SQL Server 2000 service group.

To modify the IP resource

- 1 From the Cluster Explorer configuration tree, select the IP resource in the SQL Server 2000 service group.
- 2 In the Properties View, click the **Edit** icon for the **Address** attribute.
- 3 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the virtual IP address at Site B.
 - Click **OK**.
- 4 In the Properties View, click the **Edit** icon for the **SubNetMask** attribute.
- 5 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the subnet mask at Site B.
 - Click **OK**.
- 6 From the **File** menu of Cluster Explorer, click **Close Configuration**.

Verifying the campus cluster: Switching the service group

Failover simulation is an important part of configuration testing.

To verify the campus cluster is functioning properly

- 1 Bring the service group online on one node:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Online**, and click the appropriate system from the menu.
- 2 Switch the service group to the other node:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To**, and click the appropriate system from the menu.

Setting the ForceImport attribute to 1 after a site failure

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

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

You must set the ForceImport attribute for the VMDg resource to 1 after a site failure to ensure proper failover.

To set the ForceImport attribute to 1 from the Java Console

- 1 From the Cluster Explorer configuration tree, select the VMDg resource in the SQL Server 2000 service group.
- 2 In the Properties View, click the **Edit** icon for the **ForceImport** attribute.
- 3 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system in Site B.
 - Select the **ForceImport** check box.

- Click **OK**.
- 4 From the **File** menu of Cluster Explorer, click **Close Configuration**.
 - 5 After the failover takes place, revert the ForceImport attribute to its original value.

You can also set the ForceImport attribute value using the command line. The command for implementing the force import setting in VCS is:

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

Example:

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

Import is forced on vmdg_Dg1.

Modifying the SQL 2000 service group to add VMDg and MountV resources

If you add a new SQL Server database, you must add the VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.

- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as **CRITICAL**.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

Deploying SFW HA for Campus Cluster: New SQL Server 2005 Installation

This chapter covers the following topics:

- [“Reviewing the requirements”](#) on page 459
- [“Reviewing the configuration”](#) on page 464
- [“Configuring the storage hardware and network”](#) on page 465
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 468
- [“Configuring the cluster”](#) on page 474
- [“Configuring cluster disk groups and volumes for SQL Server 2005”](#) on page 492
- [“Managing disk groups and volumes”](#) on page 501
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 503
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 507
- [“Installing and configuring SQL Server 2005 on the second node”](#) on page 510
- [“Setting the internal name of the clustered instance”](#) on page 514
- [“Creating a SQL Server user-defined database”](#) on page 517
- [“Completing configuration steps in SQL Server”](#) on page 518
- [“Configuring the SQL Server 2005 service group for VCS”](#) on page 519

- [“Modifying the IP resource in the SQL Server 2005 service group”](#) on page 525
- [“Verifying the campus cluster: Switching the service group”](#) on page 526
- [“Setting the ForceImport attribute to 1 after a site failure”](#) on page 526
- [“Modifying the SQL 2005 service group to add VMDg and MountV resources”](#) on page 527

This chapter provides information on how to install and configure a new Veritas Storage Foundation HA environment for SQL Server in a campus cluster. A campus cluster environment provides high availability and disaster recovery that extends beyond local clustering and mirroring at a single site, but is not as complex as SFW HA DR solution with replication.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for SQL Server 2005.

See [“Using the Solutions Configuration Center”](#) on page 31.

Table 11-1 on page 457 outlines the high-level objectives and the tasks to complete each objective.

Table 11-1 Task list: SQL Server 2005 Campus Cluster configuration

Objective	Tasks
“ Reviewing the requirements ” on page 459	<ul style="list-style-type: none"> Verifying hardware and software prerequisites
“ Reviewing the configuration ” on page 464	<ul style="list-style-type: none"> ■ Understanding active/passive configuration ■ Reviewing the sample configuration
“ Configuring the storage hardware and network ” on page 465	<ul style="list-style-type: none"> ■ Setting up the network and storage for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“ Installing Veritas Storage Foundation HA for Windows ” on page 468	<ul style="list-style-type: none"> ■ Verifying the driver signing options for the system ■ Installing SFW and VCS (automatic installation) and installing Veritas Cluster Server Application Agent for Microsoft SQL Server ■ Restoring driver signing options for the system
“ Configuring cluster disk groups and volumes for SQL Server 2005 ” on page 492	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the VEA
“ Configuring the cluster ” on page 474	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Running the Veritas Cluster Server Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster
“ Installing and configuring SQL Server 2005 on the first node ” on page 503	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2005 ■ Setting SQL Server services to manual start

Table 11-1 Task list: SQL Server 2005 Campus Cluster configuration

Objective	Tasks
“Preparing to install SQL Server 2005 on the second node” on page 507	<ul style="list-style-type: none"> ■ Stopping the SQL service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Mounting the volumes (adding drive letters) ■ Removing shared SQL files
“Installing and configuring SQL Server 2005 on the second node” on page 510	Installing and configuring SQL Server 2005
“Setting the internal name of the clustered instance” on page 514	Setting the internal name of the clustered instance
“Creating a SQL Server user-defined database” on page 517	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server
“Configuring the SQL Server 2005 service group for VCS” on page 519	Creating a SQL Server service group using the SQL Server Configuration Wizard
“Modifying the IP resource in the SQL Server 2005 service group” on page 525	Modifying the Address and SubNetMask attributes if the sites are in different subnets.
“Setting the ForceImport attribute to 1 after a site failure” on page 526	If a site failure occurs, setting the ForceImport attribute of the VMDg resource to 1 to ensure proper failover.

Reviewing the requirements

The campus cluster solution allows for clustered systems with mirrored or synchronously replicated storage arrays to be implemented in separate datacenters, located either within the same building or separate buildings. For example, datacenter A could be located in building A and datacenter B located in building B. This guide will refer to these different areas as Site A and Site B.

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

Disk space requirements

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

[Table 11-2](#) on page 459 estimates disk space requirements for SFW HA.

Table 11-2 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | | |
|--|---|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | ■ | Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) |
| | ■ | Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) |
| | ■ | Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ | Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) |
| | ■ | Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | ■ | Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | ■ | Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ | Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) |
| | ■ | Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |

Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required)	<ul style="list-style-type: none"> ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition
--	---

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 463.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- You must have write permissions for the Active Directory objects corresponding to all the nodes.

- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

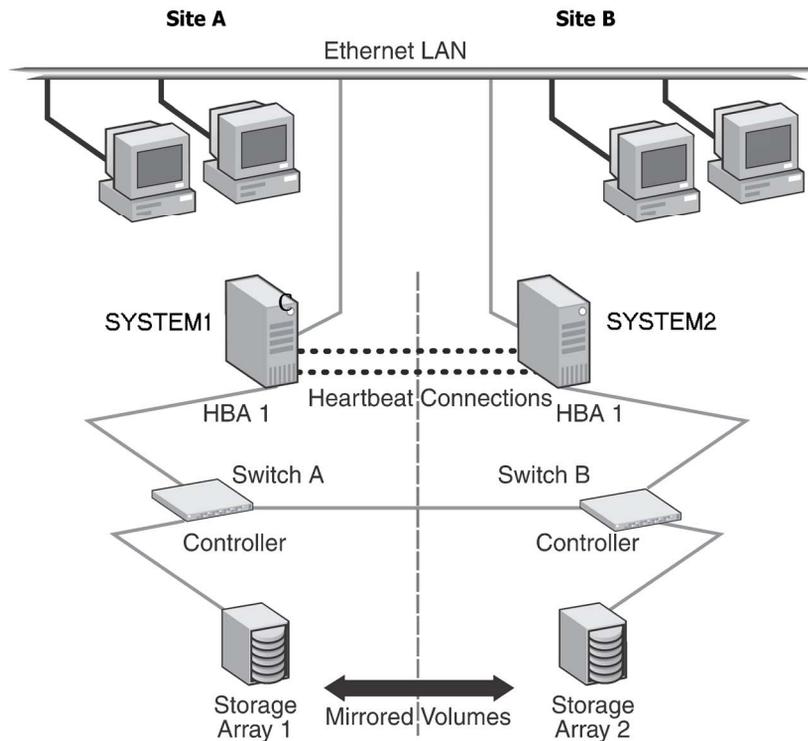
- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

This chapter uses the example of a two-node campus cluster with each node in a separate site (Site A or Site B). In this example, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array.

Figure 11-1 illustrates a Active-Passive configuration for SQL Server with one to one failover capabilities. In an active/passive configuration, the active node of the cluster hosts the SQL 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. In this case, the SQL virtual server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

Figure 11-1 SQL Server 2005 Campus Cluster: Active-Passive configuration



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 should contain the same number of disks on each site for the mirrored volumes.

Following the workflow in the Solutions Configuration Center

The Solutions Configuration Center helps you through the process of installing and configuring a new Veritas Storage Foundation HA Campus Cluster environment for one or more instances of SQL Server 2005, in either an active-passive configuration.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.

- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 11-3](#) on page 468 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 11-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 469.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

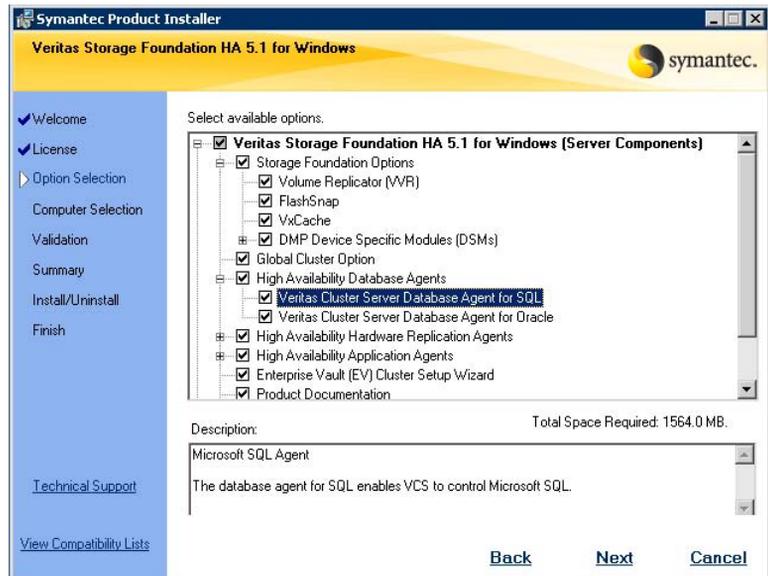
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL

Required to configure high availability for SQL Server.

Client

Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option

Required for a disaster recovery configuration only.

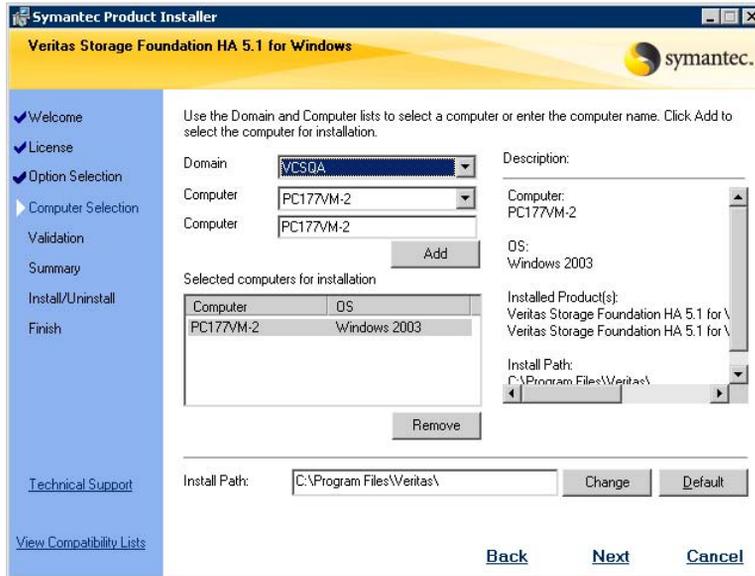
Veritas Volume Replicator

For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

High Availability Hardware Replication Agents

If you plan to use hardware replication, select the appropriate hardware replication agent.

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Set the required permissions:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.

- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

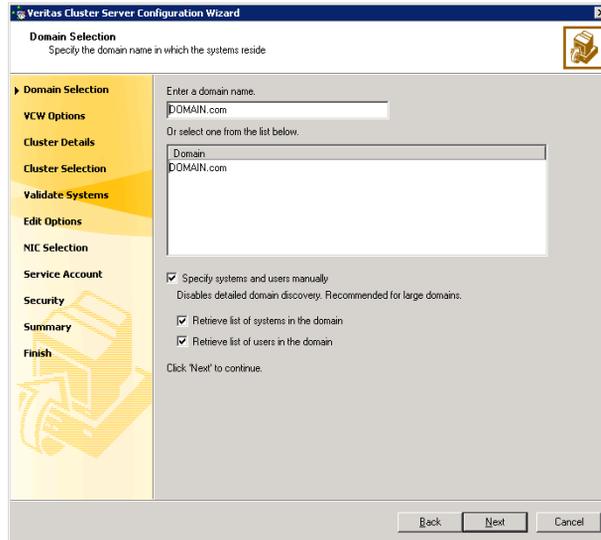
Note: If you are setting up a cluster with multiple instances of SQL, plan to add all nodes for all instances to the cluster the first time that you run the wizard. If you do that, you do not need to run the wizard again later to add the nodes.

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

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

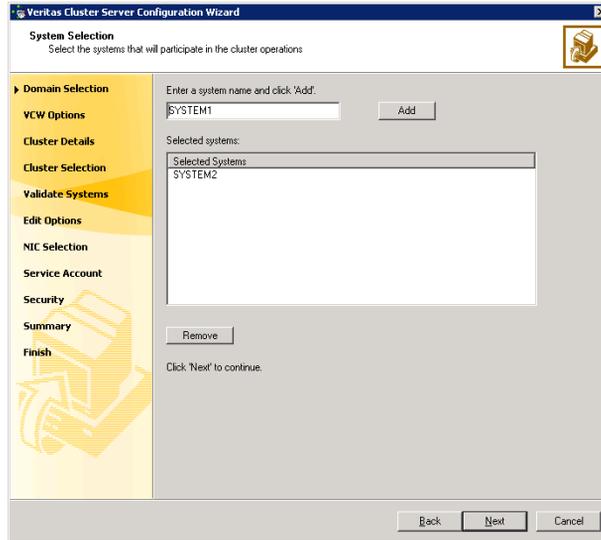
- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 8](#) on page 478.

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

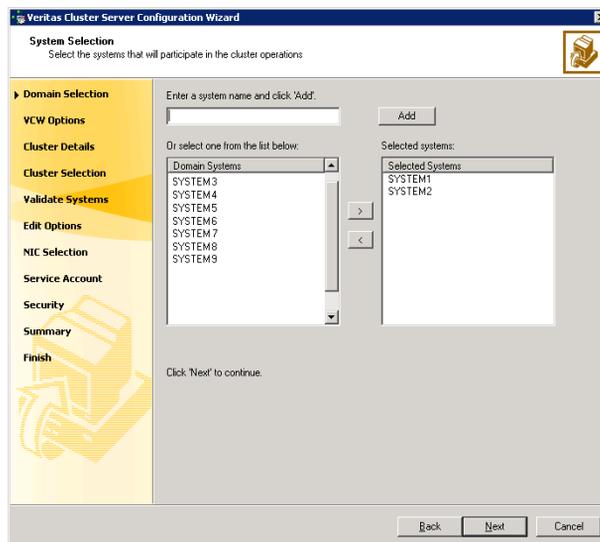
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 477. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 478.

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



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

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

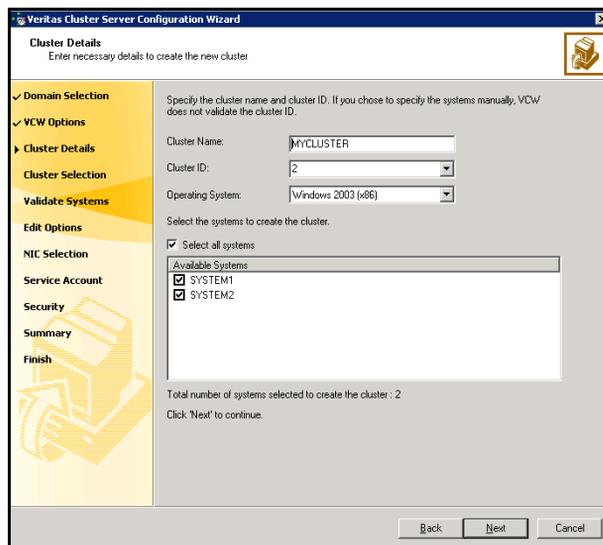
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



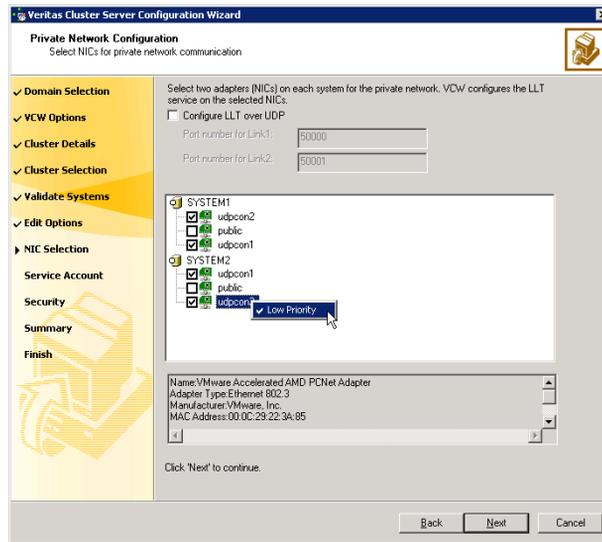
- Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
- Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

Caution: If you chose to specify systems and users manually in [step 4](#) on page 476 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 that the systems are running.
- Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

- 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](#) on page 478, proceed to the next step. Otherwise, proceed to [step 12](#) on page 482.
- 11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.
 Do one of the following:

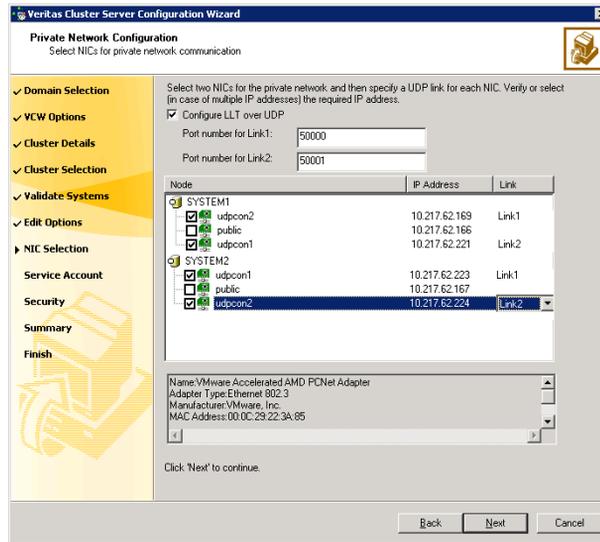
- To configure the VCS private network over Ethernet



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

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

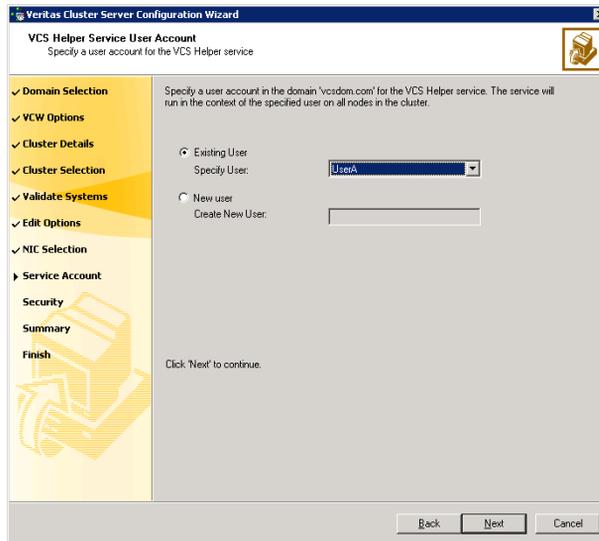
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



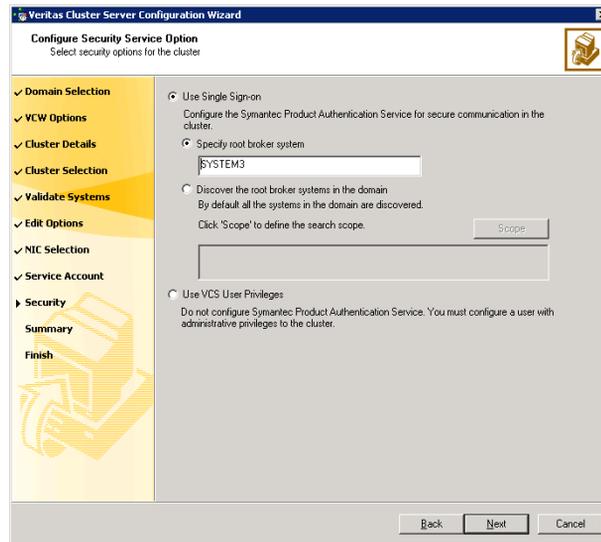
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 476, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 11-4](#) on page 484 contains some more examples of search criteria.

Table 11-4 Search criteria examples

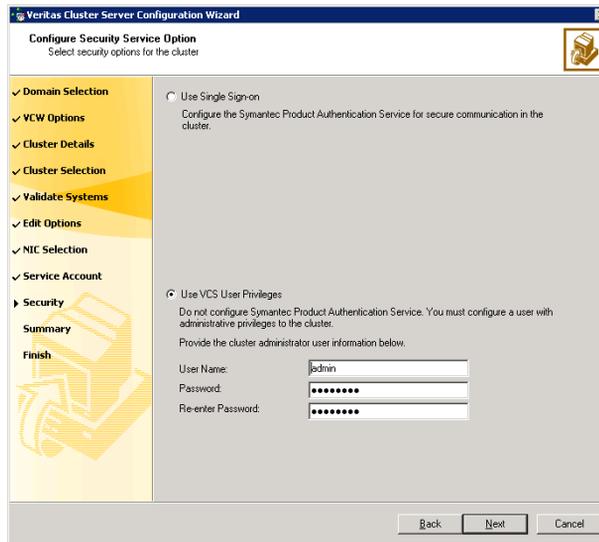
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for

the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

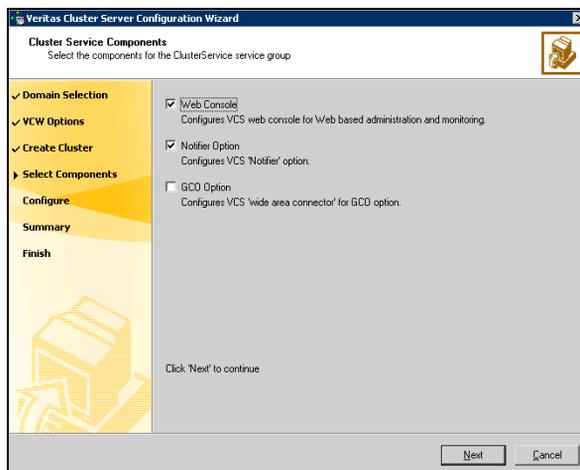
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



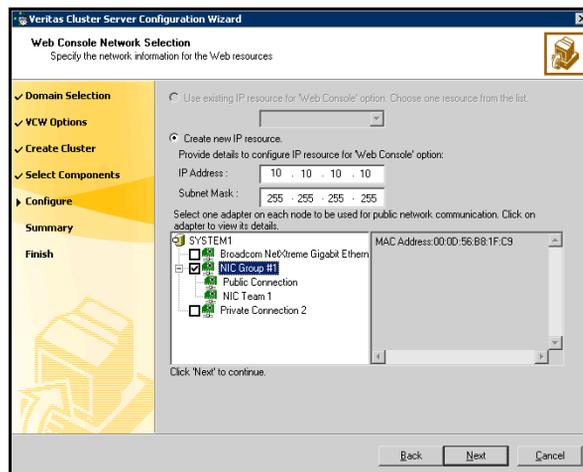
- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 487.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 488.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

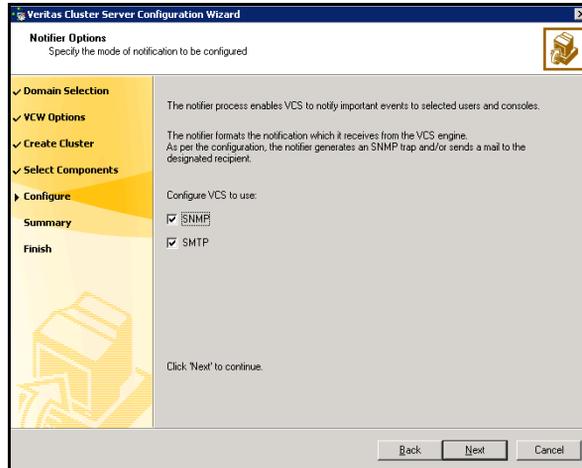
- 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 488. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

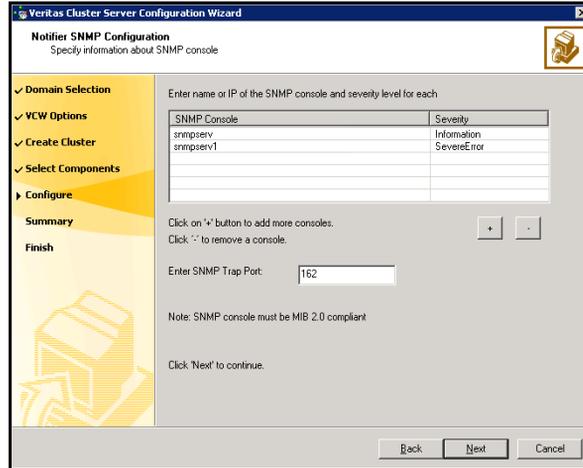
To configure notification

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



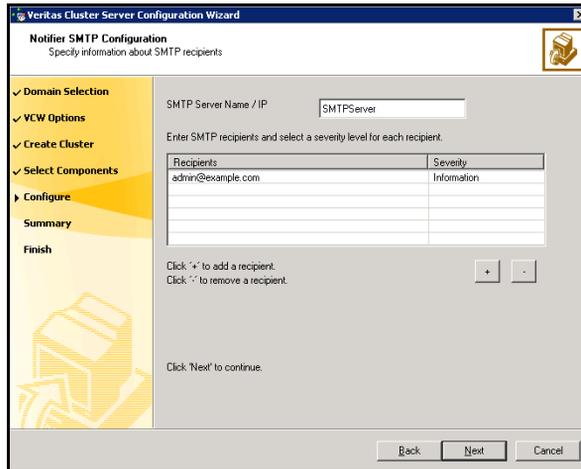
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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



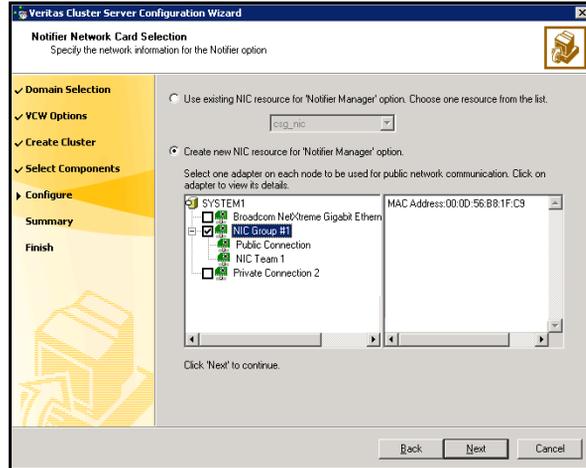
- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.



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

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



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

Configuring cluster disk groups and volumes for SQL Server 2005

Before installing SQL Server, you must create disk groups and mirrored volumes using the VEA console installed with SFW. This is also an opportunity to increase the size of existing volumes, add storage groups, and create volumes to support additional databases for storage groups.

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: Create a cluster disk group and volumes for the SQL instance on only one node of a cluster. You make the volumes accessible 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the Veritas Storage Foundation Administrator's Guide for more information.

Prerequisites for configuring cluster disk groups and volumes

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

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.

Complete the following tasks before you create the cluster disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size. Also, consider
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

Sample disk group and volume configuration

On the first node of the cluster where the SQL instance is going to be installed, you first create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL service. Create a 100 MB volume for this purpose.

You may want to place user database files in a separate cluster disk group from the system database files, for example, by creating INST1_SHARED_DG for system files and INST1_USER_DG for user database files.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

Configuring the disks and volumes

Ensure that each disk group has the same number of disks on each site. Each volume must be a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

While creating the dynamic disk groups and volumes at Site A, note carefully which disks and volumes are allocated. These will later become the Site A plexes for the mirrors.

See the following sections:

- [“Creating a dynamic \(cluster\) disk group”](#) on page 495
- [“Creating a volume”](#) on page 497

Considerations when creating new volumes

Consider the following when creating new volumes.

- For campus clusters, when creating a new volume, you must select the “mirrored across enclosures” option.
- Choosing “Mirrored” and the “mirrored across” option without having two enclosures that meet requirements causes new volume creation to fail.
- Logging can slow performance.
- Symantec recommends using either simple mirrored (concatenated) or striped mirrored 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.
- You cannot select RAID-5 for mirroring.
- Selecting “stripe across enclosures” is not recommended because then you need four enclosures, instead of two.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.

- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.
The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic (cluster) disk group

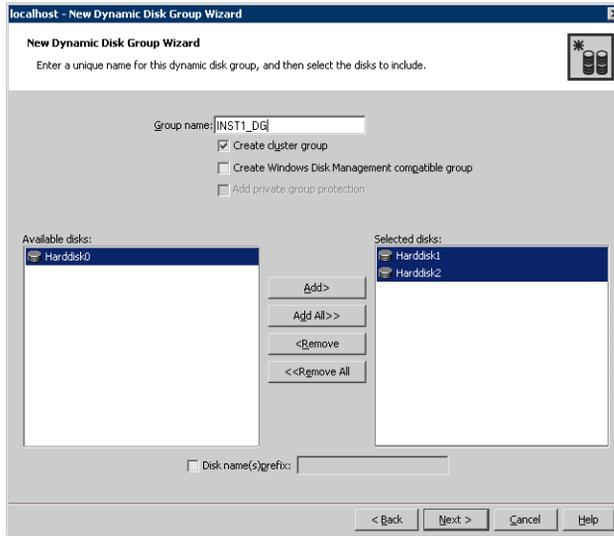
Use the following procedure to create a dynamic cluster disk group.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Proceed to create the appropriate volumes on each disk.

Creating a volume

Use this procedure to create the following volumes on the first node of the cluster:

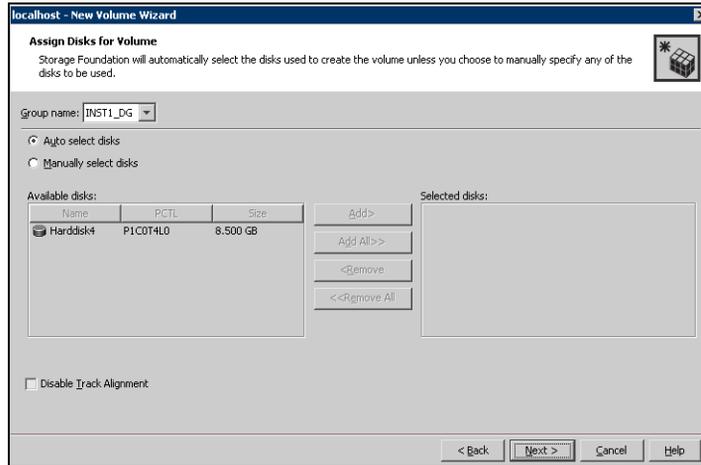
- **INST1_DATA_FILES**: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_DB1_VOL**: contains the user database files
- **INST1_DB1_LOG**: contains the user database log files
- **INST1_REGREP_VOL**: contains the list of registry keys that must be replicated among cluster systems for the SQL service.

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

To create dynamic volumes

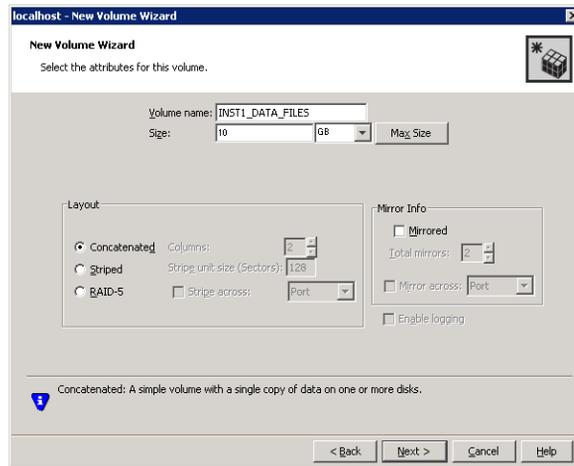
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example **INST1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume; make sure the appropriate disk group name appears in the Group name drop-down list.



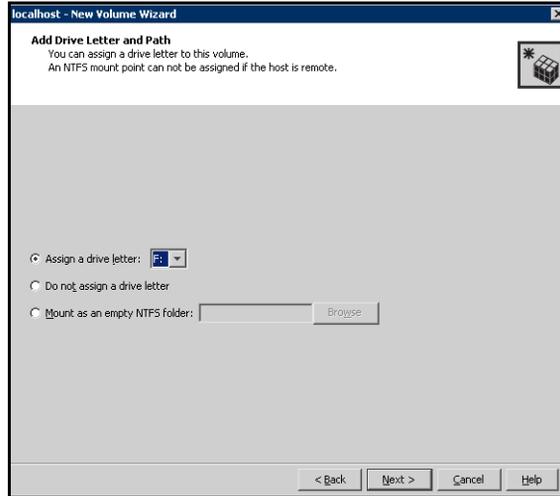
- 7 Select auto or manual disk selection and enable or disable track alignment.
 - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
 - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
 - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.
 - To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
 - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- 8 Click **Next**.

9 Specify the parameters of the volume.

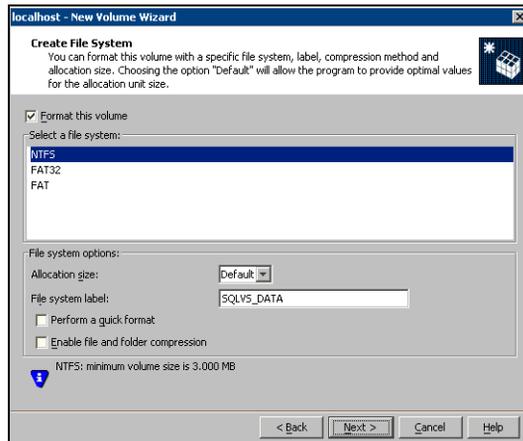


- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - Select a volume layout type. For campus clusters, select either **Concatenated** or **Striped**. Since campus clusters are mirrored volumes, you also must select the **Mirrored** checkbox.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
 - In the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
 - Verify that **Enable logging** is not selected and click **Next**.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



- 11 Click **Next**.
- 12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and select **NTFS**.
- Select an allocation file size or accept the Default.

- The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance. Click **Next**.
- 13 Click **Finish** to create the new volume.
- 14 Repeat these steps to create additional volumes.

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

Managing disk groups and volumes

This section includes the following procedures:

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

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

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.
- To move a disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - *To assign a drive letter*
Select **Assign a Drive Letter**, and select a drive letter.
 - *To mount the volume as a folder*
Select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

- 1 From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2 In the Drive Letter and Paths dialog box, click **Remove**. Click **OK** to continue.
- 3 Click **Yes** to confirm.
- 4 From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5 Click **Yes**.

Installing and configuring SQL Server 2005 on the first node

Complete the following tasks before installing SQL Server 2005:

- Verify that the cluster disk group is imported to the first node for this SQL instance
See “[Importing the cluster disk group](#)” on page 508.
- Mount the required volumes and ensure that they are assigned drive letter.
See “[Adding drive letters to mount the volumes](#)” on page 508.

Installing Microsoft SQL Server 2005

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

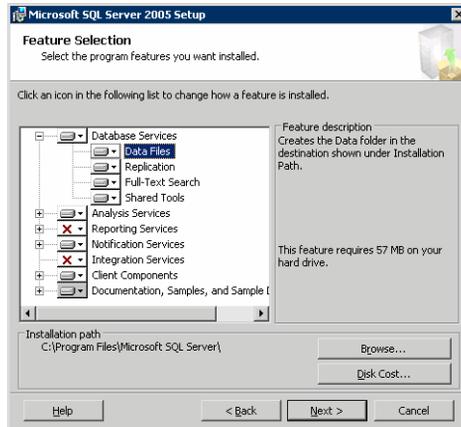
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005

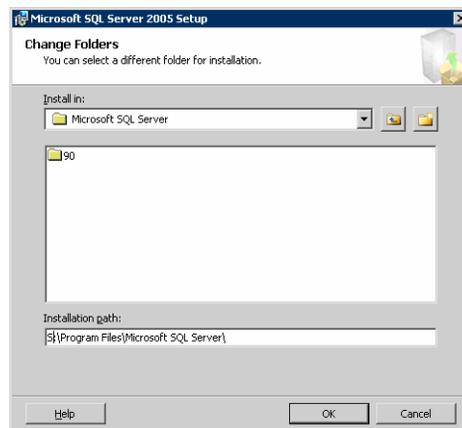
- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.

If you install optional components on one node, install the same components in the same order on other nodes.

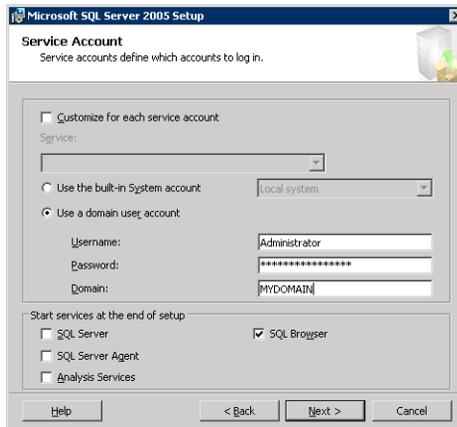
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.



- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 503, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
 Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
 - 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.

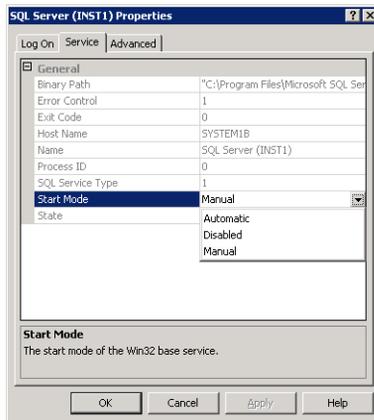
- 10 Install any SQL service packs or hotfixes if required.
- 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Setting the startup mode of the SQL Server 2005 services

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Complete the following procedures before installing SQL Server on the second or additional nodes for the SQL instance:

- “[Stopping the SQL Server 2005 service](#)” on page 507
- “[Deporting the cluster disk group](#)” on page 507
- “[Importing the cluster disk group](#)” on page 508
- “[Adding drive letters to mount the volumes](#)” on page 508
- “[Renaming shared SQL Server 2005 files](#)” on page 510

Stopping the SQL Server 2005 service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, use the Veritas Enterprise Administrator (VEA) to deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.

- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

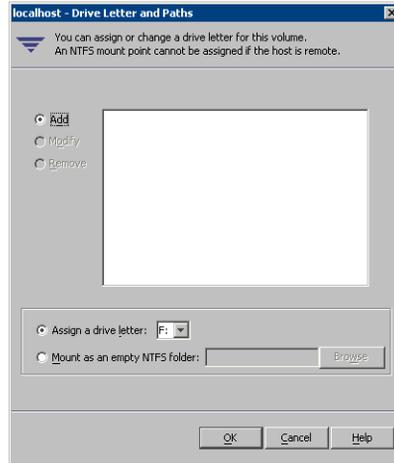
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing and configuring SQL Server 2005 on the second node

Use the following procedures to install and configure SQL Server on additional nodes for this instance:

- [“Installing SQL Server 2005 on the second node”](#) on page 510
- [“Removing shared SQL Server 2005 files”](#) on page 514

Installing SQL Server 2005 on the second node

Before installing Microsoft SQL Server 2005:

- Verify that the cluster disk group is imported to the second node
See [“Importing the cluster disk group”](#) on page 508
- Verify that the volumes are mounted (are assigned drive letters)
See [“Adding drive letters to mount the volumes”](#) on page 508

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

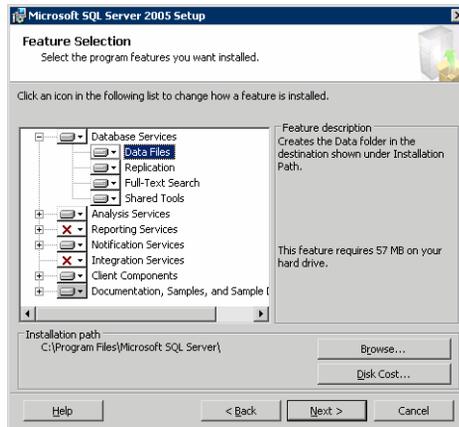
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the

installation, use the following guidelines to create an installation that will function properly in your environment.

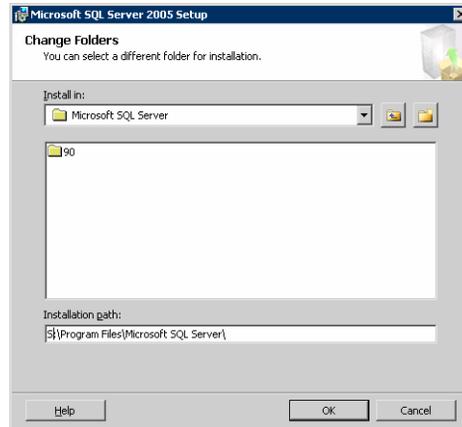
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
 Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
 If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.

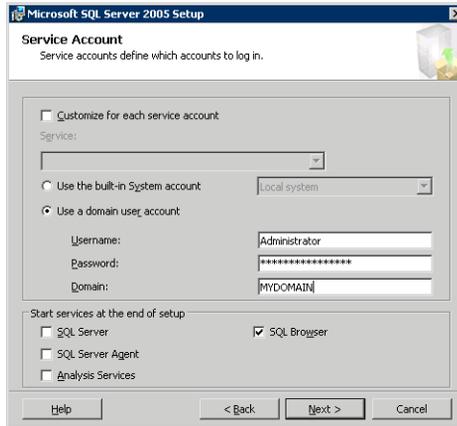


- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 511, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
 - 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
 - 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

After you complete the SQL Server installation, repeat the following procedures on the additional nodes:

- “[Preparing to install SQL Server 2005 on the second node](#)” on page 507
- “[Installing and configuring SQL Server 2005 on the second node](#)” on page 510

Removing shared SQL Server 2005 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, `INST1-VS\INST1`).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

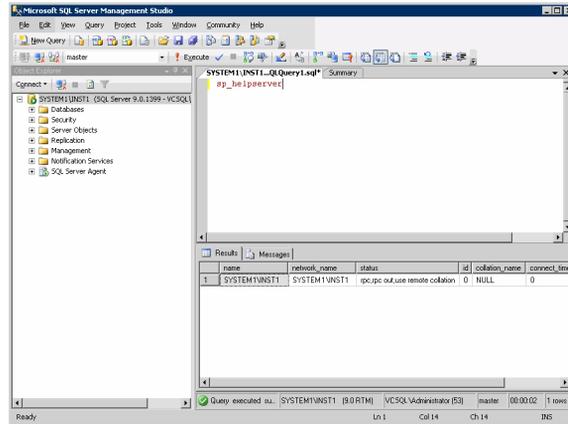
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1\INST1. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
`sp_dropserver "System_Name\Instance_Name"`
where *System_Name\Instance_Name* is the name noted in [step 3](#) on page 515.
For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press F5.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.

Creating a SQL Server user-defined database

You can use SFW HA to manage a SQL Server user-defined database. If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.

See ["Creating a volume"](#) on page 497

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Connect to the SQL Server virtual server instance for which you want to create the database (for example, `INST1-VS\INST1`).
- 3 Expand the icon associated with your server.
- 4 Right-click on **Databases** and select **New Database**.
- 5 In the New Database page, enter a name for the new database.
- 6 Click the browse button (...) in the **Path** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 7 Select and edit other file properties as desired.
- 8 Click the browse button (...) in the **Path** column for the **Transaction Log** row and browse to the location of the volume you want to create for the transaction log, and click **OK**.
- 9 Add more data files if required and configure whatever options are required for your database.
- 10 Click **OK**.

- 11 Depending on your configuration plans, you may have additional steps to complete in SQL Server.
See “[Completing configuration steps in SQL Server](#)” on page 518.
- 12 If the SQL Server service group has already been configured, you need to add the resources for the new database to the service group.
See “[Modifying the SQL 2005 service group to add VMDg and MountV resources](#)” on page 527.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Veritas Volume Replicator (VVR), Symantec recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume.

See “[Moving the tempdb database if using VVR for disaster recovery](#)” on page 518.

Moving the tempdb database if using VVR for disaster recovery

If you plan to implement a disaster recovery configuration using VVR, Symantec recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See “[Creating a volume](#)” on page 497.

Then, refer to the Microsoft Knowledge Base for the instructions on moving the tempdb database. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 224071: How to move SQL Server databases to a new location by using Detach and Attach functions in SQL Server

Refer to:

<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Configuring the SQL Server 2005 service group for VCS

Configuring the SQL Server 2005 service group involves creating an SQL Server 2005 service group and defining the attribute values for its resources. A VCS SQL Server service group is used to bring a SQL Server 2005 instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. You use the VCS SQL Configuration Wizard to configure the service group.

Prerequisites for configuring the SQL Server 2005 service group

- Verify that SFW HA, along with the VCS application agent for SQL Server 2005, is installed on all cluster nodes.
See “[Installing Veritas Storage Foundation HA for Windows](#)” on page 468.
- Verify that you have configured a VCS cluster using VCS Cluster Configuration Wizard (VCW).
See “[Configuring the cluster](#)” on page 474.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group.
See “[Importing the cluster disk group](#)” on page 508.
See “[Adding drive letters to mount the volumes](#)” on page 508.
- Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name when configuring the service group.
See “[Setting the internal name of the clustered instance](#)” on page 514.
- Assign a unique virtual IP address for the SQL Server 2005 instance. You specify this IP address when configuring the service group.

- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists on all the cluster nodes. A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2005\sample_script.sql`. Detailed monitoring is often not necessary.
- Stop the SQL 2005 Server service for the SQL instance.
See “[Stopping the SQL Server 2005 service](#)” on page 507.

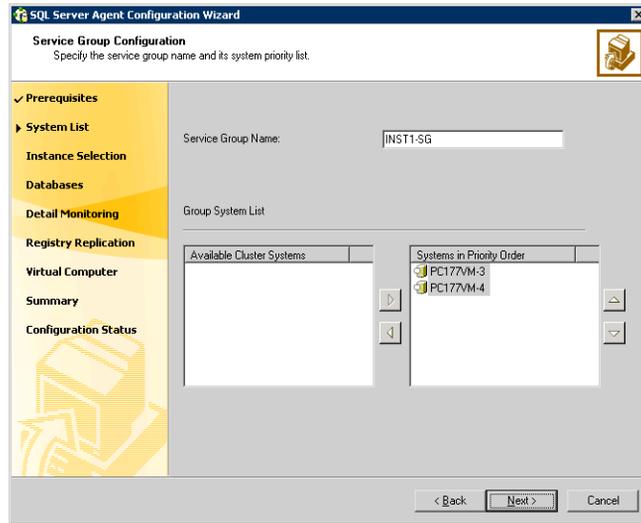
Creating the SQL Server 2005 service group

The VCS SQL Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

To create a SQL Server service group on the cluster

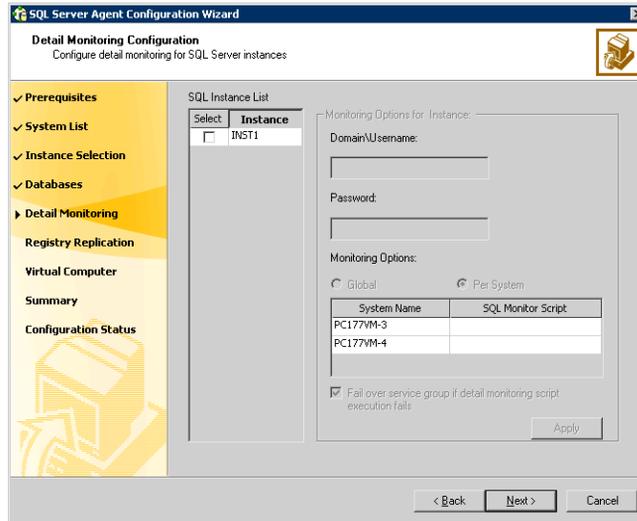
- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.

5 Specify the service group name and system list:



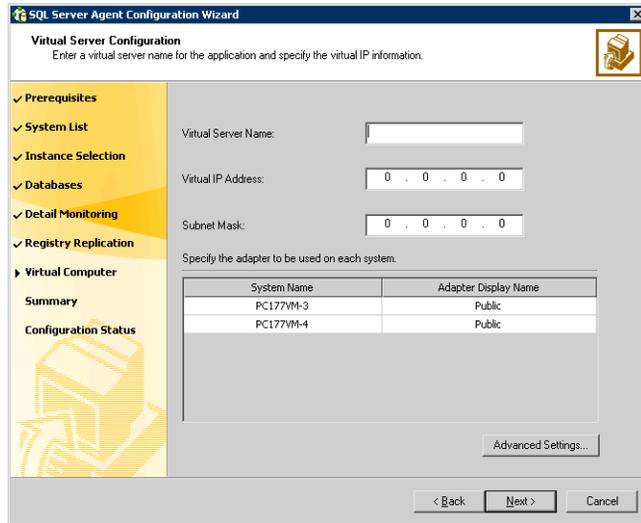
- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
 Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 521 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 522 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.

- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example `INST1-VS`. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
 The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.

- Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

If you have created a new SQL Server database, you must add VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

See “[Modifying the SQL 2005 service group to add VMDg and MountV resources](#)” on page 527.

Modifying the IP resource in the SQL Server 2005 service group

Note: This procedure is only applicable to a campus cluster with sites in different subnets.

Use the Java Console to modify the Address and SubNetMask attributes of the IP resource in the SQL Server 2005 service group.

To modify the IP resource

- 1 From the Cluster Explorer configuration tree, select the IP resource in the SQL Server 2005 service group.
- 2 In the Properties View, click the **Edit** icon for the **Address** attribute.
- 3 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the virtual IP address at Site B.
 - Click **OK**.
- 4 In the Properties View, click the **Edit** icon for the **SubNetMask** attribute.
- 5 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the subnet mask at Site B.
 - Click **OK**.
- 6 From the **File** menu of Cluster Explorer, click **Close Configuration**.

Verifying the campus cluster: Switching the service group

Failover simulation is an important part of configuration testing.

To verify the campus cluster is functioning properly

- 1 Bring the service group online on one node:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Online**, and click the appropriate system from the menu.
- 2 Switch the service group to the other node:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To**, and click the appropriate system from the menu.

Setting the ForceImport attribute to 1 after a site failure

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

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

You must set the ForceImport attribute for the VMDg resource to 1 after a site failure to ensure proper failover.

To set the ForceImport attribute to 1 from the Java Console

- 1 From the Cluster Explorer configuration tree, select the VMDg resource in the SQL Server 2005 service group.
- 2 In the Properties View, click the **Edit** icon for the **ForceImport** attribute.
- 3 In the Edit Attribute dialog box:
 - Select the **Per System** option.
 - Select the system in Site B.
 - Select the **ForceImport** check box.

- Click **OK**.

- 4 From the **File** menu of Cluster Explorer, click **Close Configuration**.
- 5 After the failover takes place, revert the ForceImport attribute to its original value.

You can also set the ForceImport attribute value using the command line. The command for implementing the force import setting in VCS is:

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

Example:

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

Import is forced on vmdg_Dg1.

Modifying the SQL 2005 service group to add VMDg and MountV resources

If you add a new SQL Server database, you must add the VMDg and MountV resources to the SQL Server service group, using the SQL Server Configuration Wizard.

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.

- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as `CRITICAL`.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

Replicated Data Clusters

This section includes the following chapters:

- [About Replicated Data Clusters](#)
- [Configuring Replicated Data Clusters for SQL 2000](#)
- [Configuring Replicated Data Clusters for SQL 2005](#)

About Replicated Data Clusters

This chapter includes the following topics:

- [“About Replicated Data Clusters”](#) on page 531
- [“How VCS Replicated Data Clusters work”](#) on page 533
- [“Setting up a Replicated Data Cluster configuration”](#) on page 534
- [“Migrating the service group”](#) on page 537

About Replicated Data Clusters

A Replicated Data Cluster (RDC) uses data replication, instead of shared storage, to assure data access to all the nodes in a cluster.

The Replicated Data Cluster configuration provides both local high availability and disaster recovery functionality in a single VCS cluster. You can set up RDC in a VCS environment using Veritas Volume Replicator (VVR.)

An RDC exists within a single VCS cluster with a primary zone and a secondary zone, which can stretch over two buildings or data centers connected with Ethernet. In an RDC configuration, if an application or a system fails, the application is failed over to another system within the current primary zone. If the entire primary zone fails, the application is migrated to a system in the secondary zone (which then becomes the new primary).

For VVR replication to occur, the disk groups containing the Replicated Volume Group (RVG) must be imported at the primary and secondary zones. The replication service group must be online at both zones simultaneously, and must be configured as a hybrid VCS service group.

The SQL Server service group is configured as a failover service group. The SQL Server service group must be configured with an online local hard dependency on the replication service group.

Note: VVR supports multiple replication secondary targets for any given primary. However, RDC for VCS supports only one replication secondary for a primary.

An RDC configuration is appropriate in situations where dual dedicated LLT links are available between the primary zone and the secondary zone but lacks shared storage or SAN interconnect between the primary and secondary data centers. In an RDC, data replication technology is employed to provide node access to data in a remote zone.

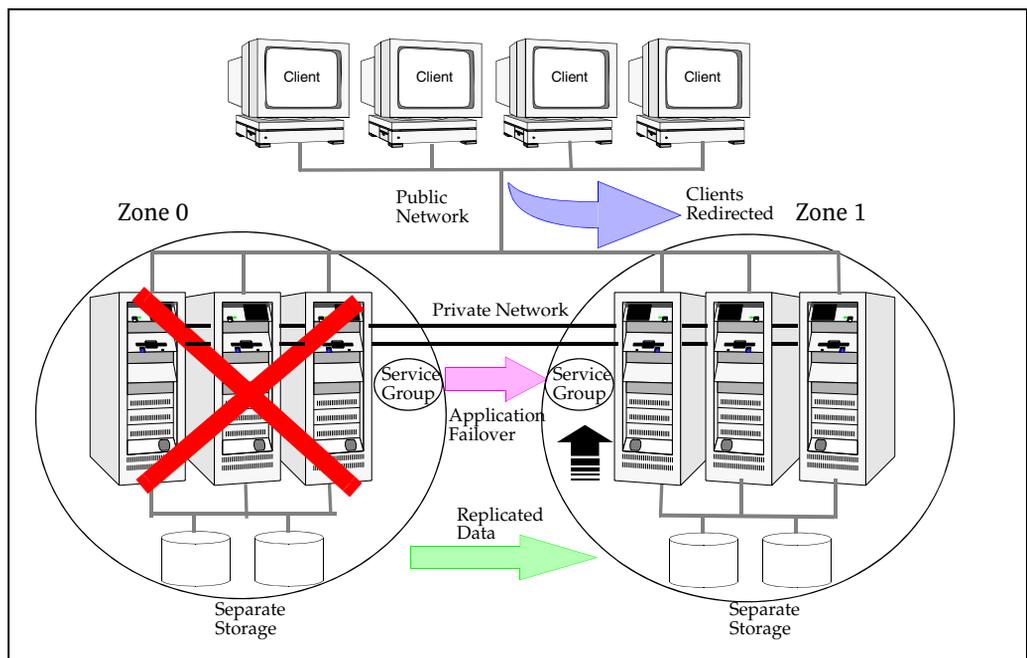
You must use dual dedicated LLT links between the replicated nodes.

How VCS Replicated Data Clusters work

To understand how a RDC configuration works, let us take the example of SQL 2000 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

SQL 2000 is installed and configured on all nodes in the cluster. The SQL 2000 data is located on shared disks within each RDC zone and is replicated across RDC zones to ensure data concurrency. The SQL Server service group is online on a system in the current primary zone and is configured to fail over in the cluster.



In the event of a system or SQL 2000 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone (zone 1). VCS also redirects clients once the application is online on the new location.

Setting up a Replicated Data Cluster configuration

In the example, SQL Server 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. In the event of a failure on the primary node, VCS can fail over the SQL 2000 instance to the second node in the primary zone.

The process involves the following steps:

- [Setting up replication](#)
- [Configuring the service groups](#)

Setting up replication

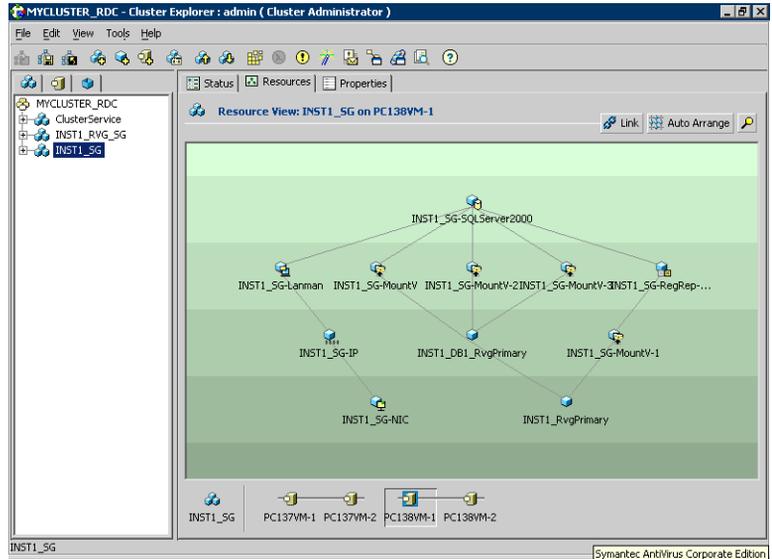
Set up replication between the shared disk groups. Use VVR to group the shared data volumes into a Replicated Volume Group, and creating the VVR Secondary on hosts in your secondary zone.

Create a Replicated Data Set (RDS) with the Primary RVG consisting of the shared volumes between the nodes in the first zone and Secondary RVG consisting of shared volumes between nodes in the second zone. Therefore, use the same Disk Group and RVG name in both zones so that the MountV resources will mount the same block devices.

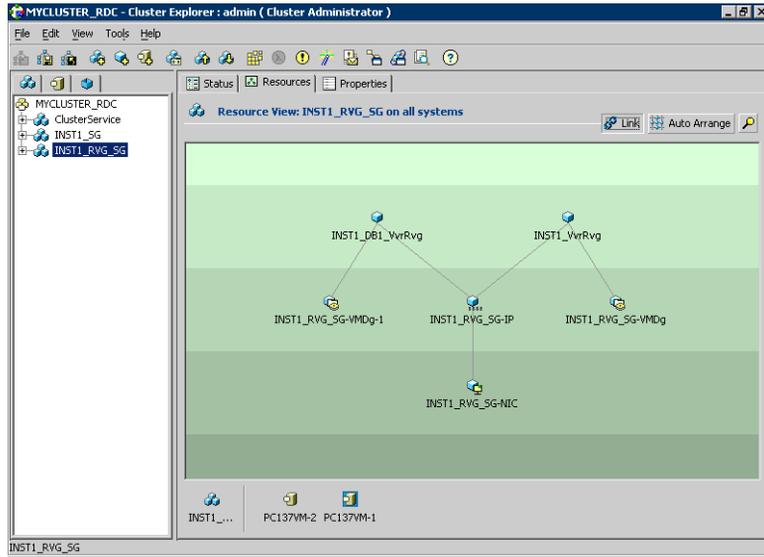
Configuring the service groups

For a successful wide-area failover, the mount points and applications must fail over to the secondary RDC zone. Additionally, the VVR secondary disk group and RVG must be imported and started on the secondary RDC zone.

The following screen from the VCS Cluster Manager (Java Console) depicts a typical SQL Server service group RDC configuration:



The following screen from the VCS Cluster Manager (Java Console) depicts a typical SQL replication service group (RVG) configuration:



Migrating the service group

In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the shared storage. In this situation, none of the nodes in the primary zone see any device.

The service group cannot fail over locally within the primary RDC zone, because the shared volumes cannot be mounted on any node. So, the service group must fail over, to a node in the current secondary RDC zone.

The RVGPrimary agent ensures that VVR volumes are made writable. The application can be started at the secondary RDC zone and run there until the problem with the local storage is corrected.

If the storage problem is corrected, you can switch the application back to the primary zone using VCS.

To switch the service group

Before switching 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.

In the **Service Groups** tab of the of the Cluster Explorer configuration tree, right-click the service group. Click **Switch To** and select the system in the primary RDC zone to switch to and click OK.

Configuring Replicated Data Clusters for SQL 2000

This chapter includes the following topics:

- [“Tasks for configuring Replicated Data Clusters for SQL 2000”](#) on page 540
- [“Reviewing the prerequisites”](#) on page 543
- [“Reviewing the configuration”](#) on page 547
- [“Configuring the storage hardware and network”](#) on page 549
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 552
- [“Configuring the cluster”](#) on page 560
- [“Configuring cluster disk groups and volumes”](#) on page 577
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 586
- [“Preparing to install SQL Server on the second node”](#) on page 589
- [“Installing SQL Server 2000 on the second node”](#) on page 593
- [“Setting the internal name of the clustered instance”](#) on page 597
- [“Configuring the VCS SQL Server service group”](#) on page 599
- [“Creating the primary system zone”](#) on page 604
- [“Registering the virtual server in the SQL Server”](#) on page 605
- [“Creating a SQL Server user-defined database”](#) on page 607
- [“Verifying the installation in the primary zone”](#) on page 609
- [“Creating a parallel environment in the secondary zone”](#) on page 610

- [“Adding the systems in the secondary zone to the cluster”](#) on page 611
- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 619
- [“Configuring a hybrid RVG service group for replication”](#) on page 630
- [“Setting a dependency between the service groups”](#) on page 645
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 646
- [“Verifying the RDC configuration”](#) on page 653
- [“Additional instructions for GCO disaster recovery”](#) on page 654

Tasks for configuring Replicated Data Clusters for SQL 2000

Configure the high availability and SQL components on the primary and secondary zones. Then complete the Replicated Data Set solution by configuring the components for both zones.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table 13-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Reviewing the prerequisites” on page 543	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 547	<ul style="list-style-type: none">■ Understanding active-passive configuration and zone failover in a RDC environment■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 549	<ul style="list-style-type: none">■ Setting up the storage hardware for a cluster environment■ Verifying the DNS entries for the systems on which SQL will be installed

Table 13-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 552	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ■ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ■ Configuring VxSAS
“Configuring the cluster” on page 560	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 577	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2000 on the first node” on page 586	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Configuring SQL services
“Preparing to install SQL Server on the second node” on page 589	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters ■ Removing shared SQL files from the cluster disk group

Table 13-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Installing SQL Server 2000 on the second node” on page 593	Installing and configuring SQL Server 2000
“Setting the internal name of the clustered instance” on page 597	Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 599	Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating the primary system zone” on page 604	<ul style="list-style-type: none"> ■ Creating the primary system zone ■ Adding the nodes to the primary zone
“Registering the virtual server in the SQL Server” on page 605	Registering the virtual server in the SQL Server
“Creating a SQL Server user-defined database” on page 607	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server ■ Adding resources for a user-defined database in VCS
“Verifying the installation in the primary zone” on page 609	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“Creating a parallel environment in the secondary zone” on page 610	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Configuring disk groups and volumes for SQL
“Setting up the Replicated Data Sets (RDS)” on page 619	Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones

Table 13-1 Tasks for configuring Replicated Data Clusters for SQL 2000

Objective	Tasks
“Configuring a hybrid RVG service group for replication” on page 630	<ul style="list-style-type: none"> ■ Creating a hybrid Replicated Volume Group (RVG) service group ■ Configuring the hybrid RVG service group
“Setting a dependency between the service groups” on page 645	Setting up a dependency from the VVR RVG Service Group to the SQL Server Service Group
“Adding the nodes from the secondary zone to the RDC” on page 646	<ul style="list-style-type: none"> ■ Adding the nodes from the secondary zone to the RVG service group ■ Configuring the IP resources for failover ■ Adding the nodes from the secondary zone to the SQL Server service group
“Verifying the RDC configuration” on page 653	Verifying that failover occurs first within zones and then from the primary to the secondary zone

Reviewing the prerequisites

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation. This replication recovery solution requires installation and configuration at a primary zone and a secondary zone.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at the following URL:

<http://www.symantec.com/business/support/index.jsp>

Supported software

The following software is supported:

- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL and the Veritas Volume Replicator
- For a Microsoft SQL Server 2000 environment, any of the following SQL Servers and their operating systems:

Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)	<ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions)
Microsoft SQL Server 2000 (64-bit) Enterprise Edition	<ul style="list-style-type: none">■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)
Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)	<ul style="list-style-type: none">■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions)

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Shared disks to support applications that migrate between nodes in the cluster. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See “[Best practices](#)” on page 546.
- Memory: minimum 1 GB of RAM per server for SFW HA.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL Virtual Server
 - A minimum of one static IP address for each physical node in the cluster
 - For VVR, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.

- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *VCS Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on `C:\WINDOWS` of one node, installations on all other nodes must be on `C:\WINDOWS`. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- When installing, install only in a single domain.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).

When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.

- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Note: Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `xclus UseSystemBus ON` command.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 13-2](#) on page 547 estimates disk space requirements for SFW HA.

Table 13-2 Disk space requirements

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

Reviewing the configuration

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the 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

You should have these IP addresses available before you start deploying your environment.

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:

Primary zone

SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
INST1_SG	Microsoft SQL Server 2000 service group
INST1-VS	Virtual SQL Server cluster
INST1	SQL Instance Name
INST1_DG	Cluster disk group for SQL system database and files
INST1_DATA_FILES	Volume for Microsoft SQL Server system data files
INST1_REGREP_VOL	Volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	Replicator log volume required by VVR
INST1_DB1_DG	Cluster disk group for SQL Server user-defined database and files
INST1_DB1_VOL	Volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	Volume for storing a Microsoft SQL Server user-defined database log file
INST1_DB1_REPLOG	Replicator log volume required by VVR for SQL user-defined database

Secondary zone

SYSTEM3 & SYSTEM4 First and second nodes of the secondary zone

All the other parameters are the same as on the primary zone.

RDS and VVR Components

INST1_RDS	RDS name for SQL system database and files
INST1_RVG	RVG name for SQL system database and files

INST1_RVG_SG	Replication service group for SQL system database and files
INST1_DB1_RDS	RDS name for SQL Server user-defined database and files
INST1_DB1_RVG	RVG name for SQL Server user-defined database and files
INST1_DB1_RVG_SG	Replication service group for SQL Server user-defined database and files

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.

- 4 Ensure the public network adapter is the first bound adapter:
 - From the **Advanced** menu in the Network Connections window, click **Advanced Settings**.
 - In the **Adapters and Bindings** tab, verify the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the **Public status** dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the **Public Status** dialog box, on the **General** tab, click **Properties**.
- 7 In the **Public Properties** dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the options to install VVR and the Veritas Cluster Server Database Agent for SQL. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 13-3](#) on page 552 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 13-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 553.

To change the driver signing options on each local system

- 1 Open the Control Panel and click **System**.
- 2 Click the Hardware tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.

If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

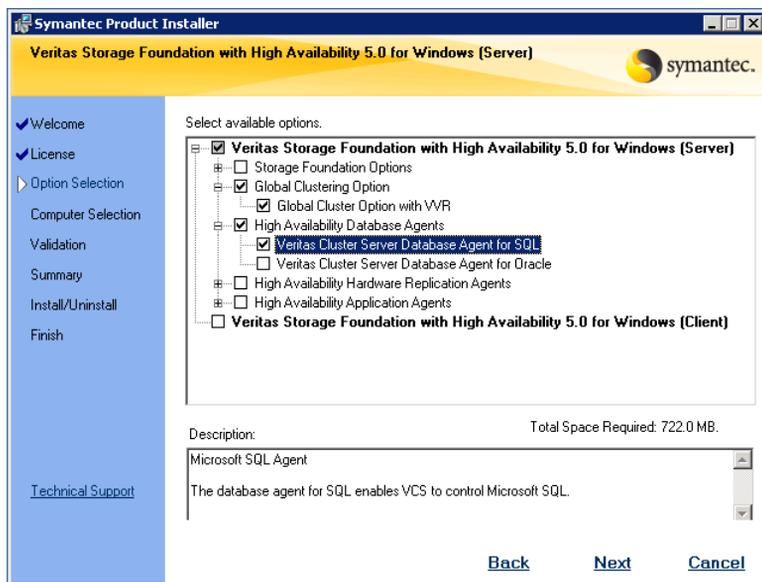
To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation HA 5.1 for Windows**.
- 4 Do one of the following:



- Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.
 - 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.
 - 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.

- 8 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.
- 9 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:

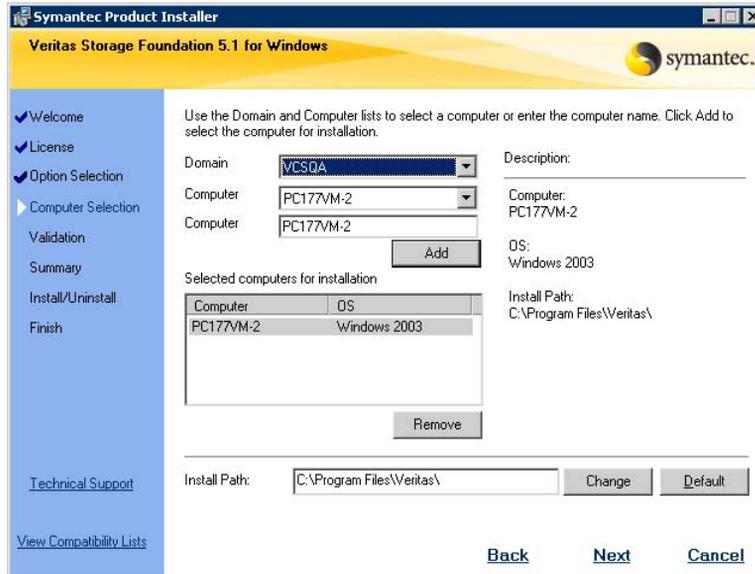


Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

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

Veritas Volume Replicator To use VVR for replication, you must select the option to install VVR.

10 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Configuring VxSAS

You can run the VVR Service Configuration (VxSAS) wizard after you install SFW HA on both the primary and secondary nodes. When you run the wizard, you can then specify the primary and secondary sites in one step.

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.
Read the information provided on the Welcome page and click **Next**.
- 2 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 the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts. Click **Next**.

- 3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

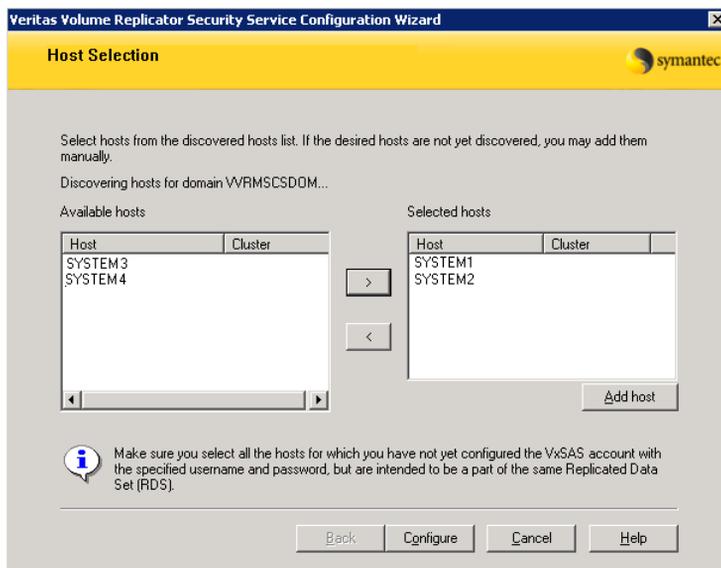
Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

- 4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

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

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.

- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

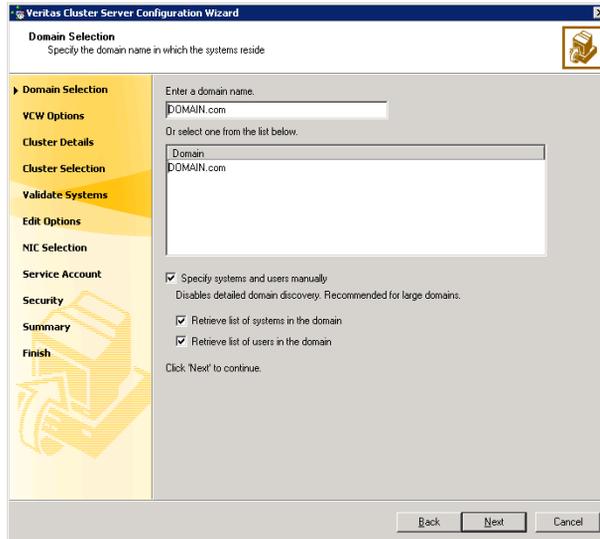
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Note: Add only systems in the primary zone (zone 0) to the cluster at this time.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

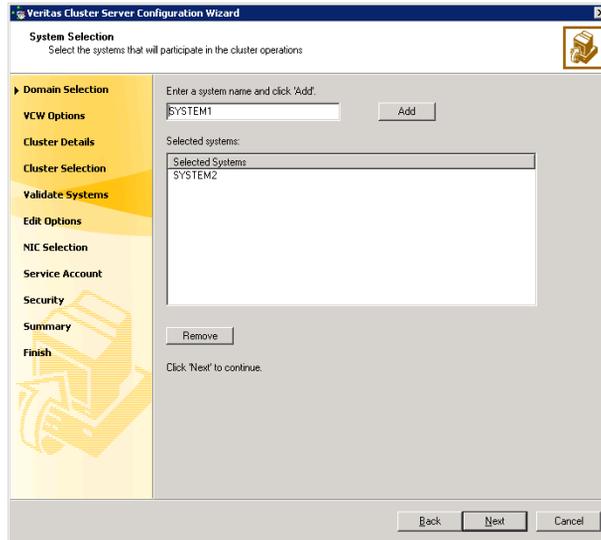
- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 8](#) on page 564.

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

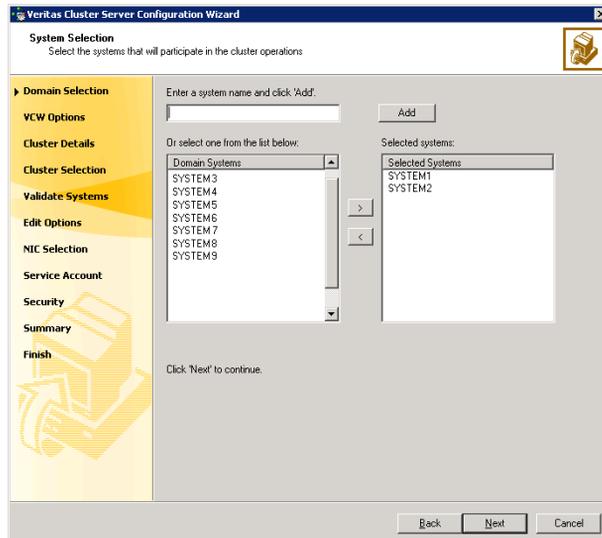
- Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**. If you chose to retrieve the list of systems, proceed to [step 6](#) on page 563. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 564.

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



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

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

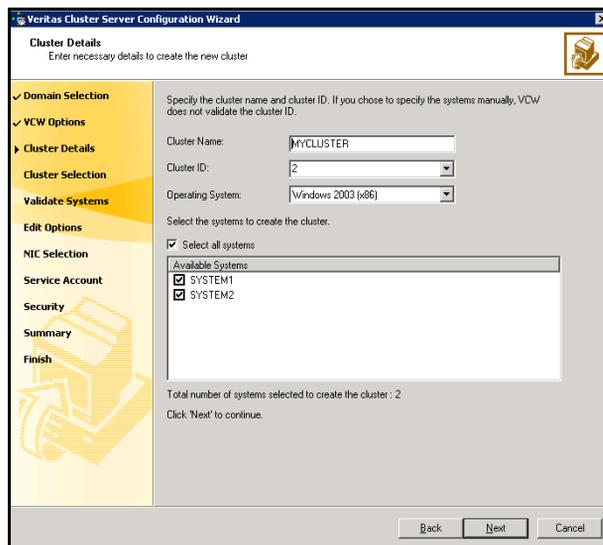
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 562 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

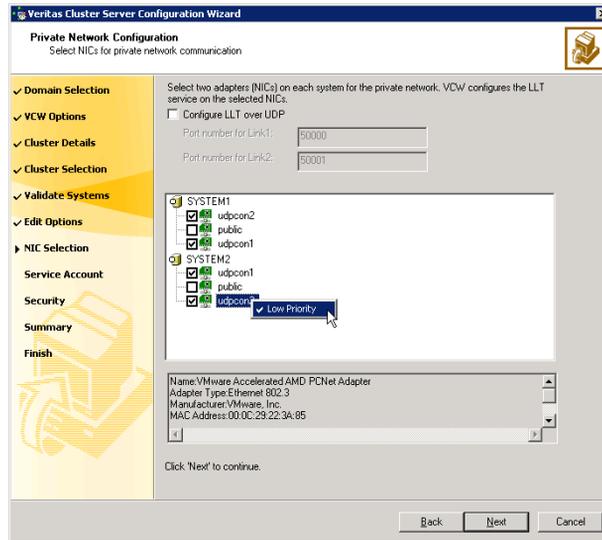
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 564, proceed to the next step. Otherwise, proceed to [step 12](#) on page 568.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

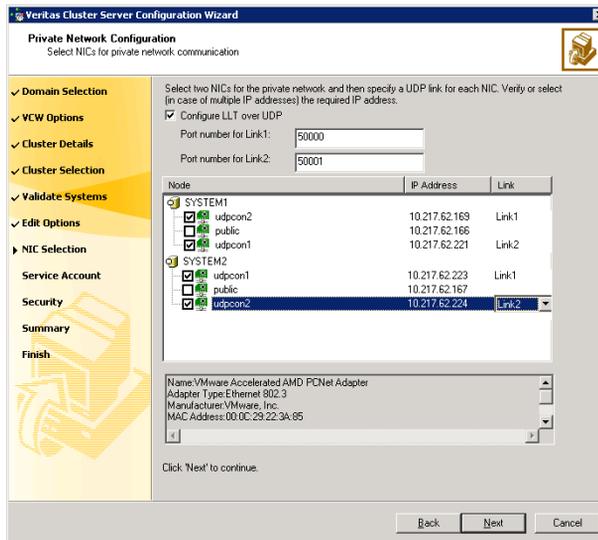
- To configure the VCS private network over Ethernet



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

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

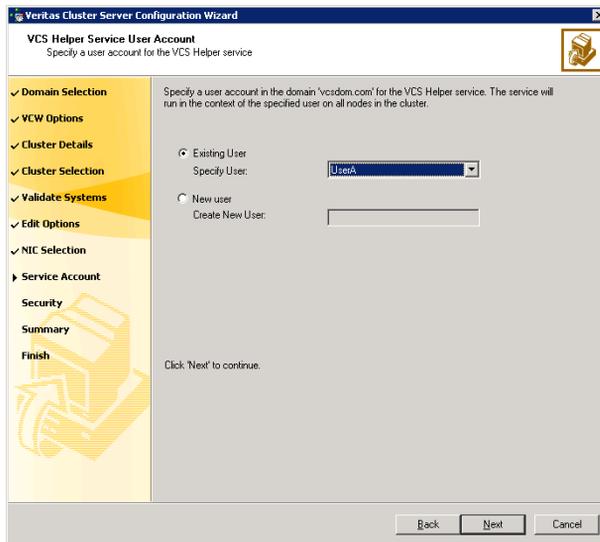
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 562, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

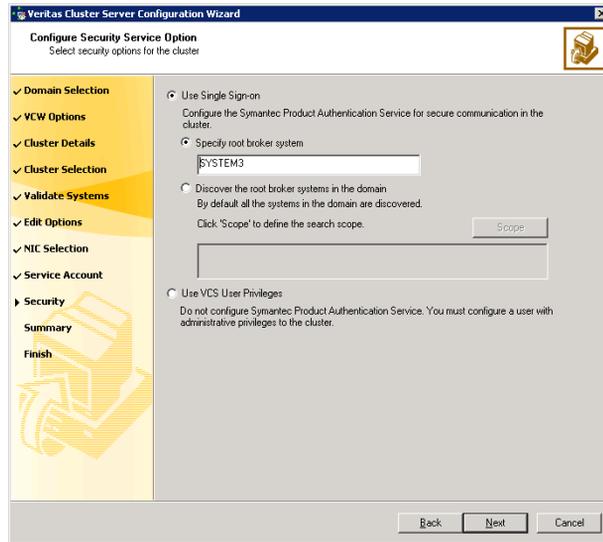
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

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

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 13-4](#) on page 570 contains some more examples of search criteria.

Table 13-4 Search criteria examples

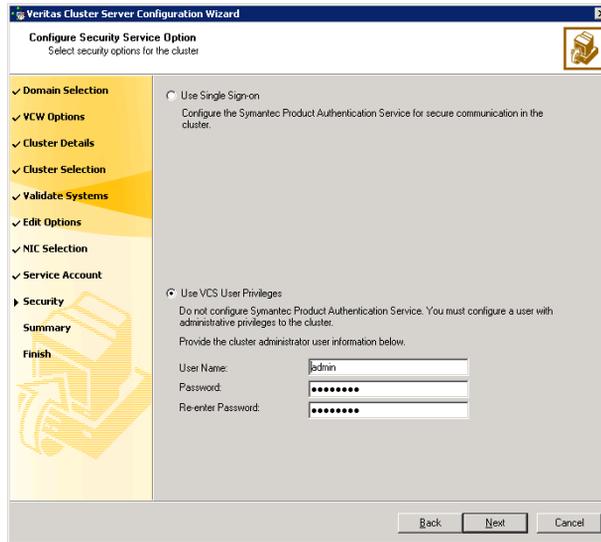
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

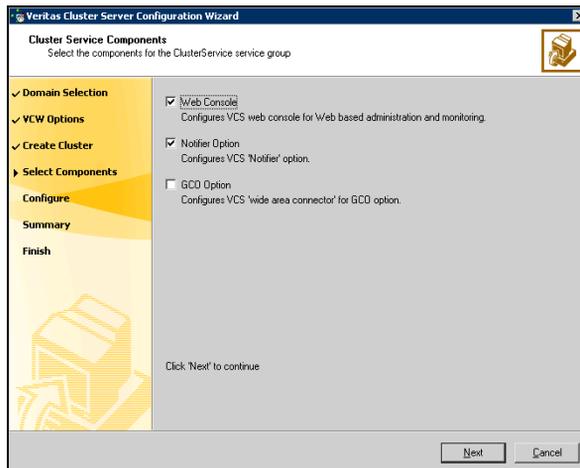
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 573.

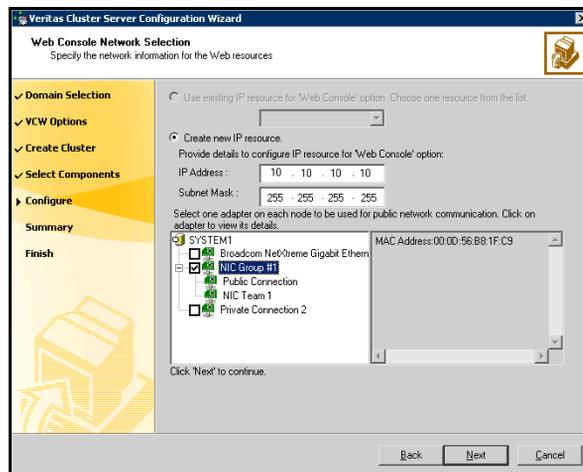
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 574.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



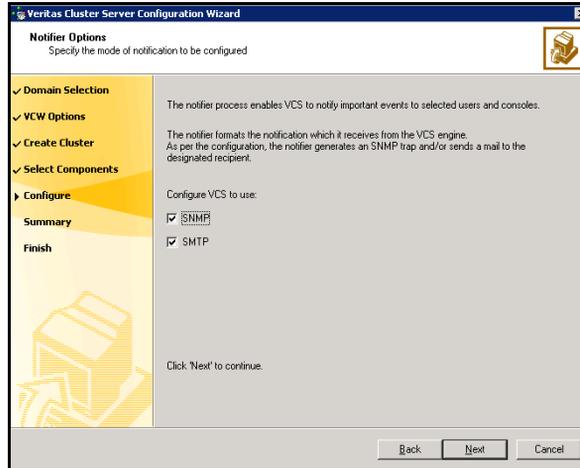
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 574. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

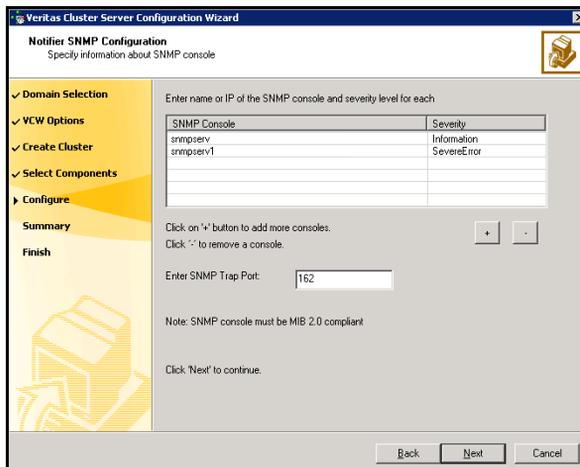
To configure notification

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

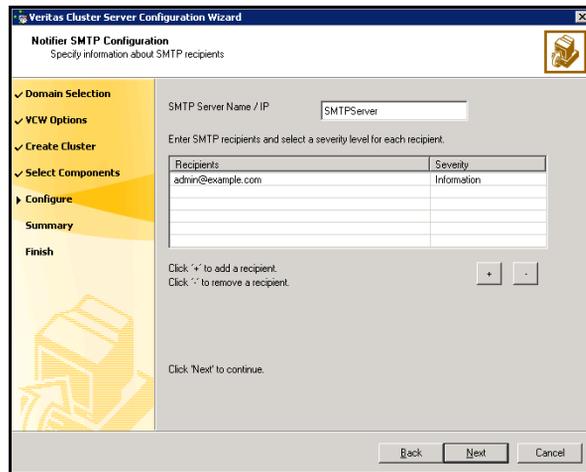


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

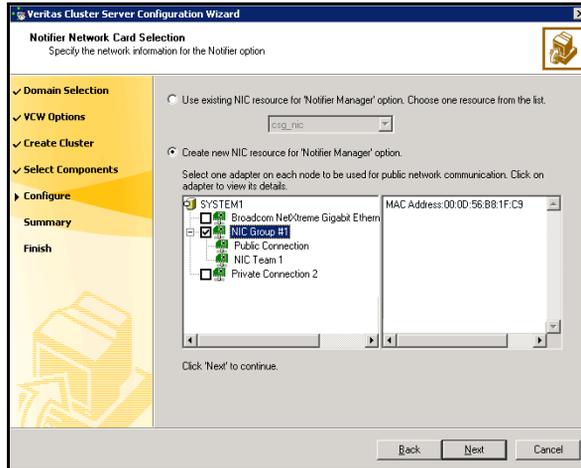


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

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

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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 disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the disk group are shared and are available from all nodes in each zone. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_REPLOG** contains the VVR Replicator Log.
You can create this volume later while setting up the Replicated Data Sets. See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 619.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and files.

The following disk group and volumes may be created now or later in the configuration process.

- **INST1_DB1_DG** is the disk group for the SQL Server user-defined database and files
- **INST1_DB1_VOL** contains the user database files
- **INST1_DB1_LOG** contains the user database log files
- **INST1_DB1_REPLOG** contains the VVR Storage Replicator Log.

Warning: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Creating a cluster disk group

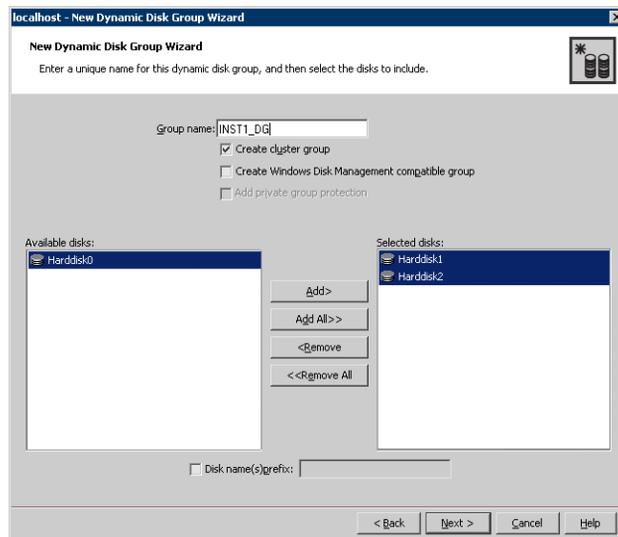
Create a cluster disk group on the first node of the cluster. You can repeat the process to create the disk group for the SQL Server user-defined database and files at this time or later.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
 Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes in the cluster disk group (INST1_DG) for the system files on the first node of the cluster:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- INST1_REPLOG contains the VVR Storage Replicator Log.
You can create the INST1_REPLOG volume at this time or during the process of creating Replicated Data Sets.
See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 619.

You can create the INST1_DB1_VOL and INST1_DB1_LOG volumes at this time or when you create a SQL user-defined database.

See “[Creating a SQL Server user-defined database](#)” on page 607.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and the log files. The following volumes in the disk group (INST1_DB1_DG) for the SQL Server user-defined database and files may be created now or later in the configuration process:

- INST1_DB1_VOL contains the user database files
- INST1_DB1_LOG contains the user database log files
- INST1_DB1_REPLOG contains the VVR Storage Replicator Log

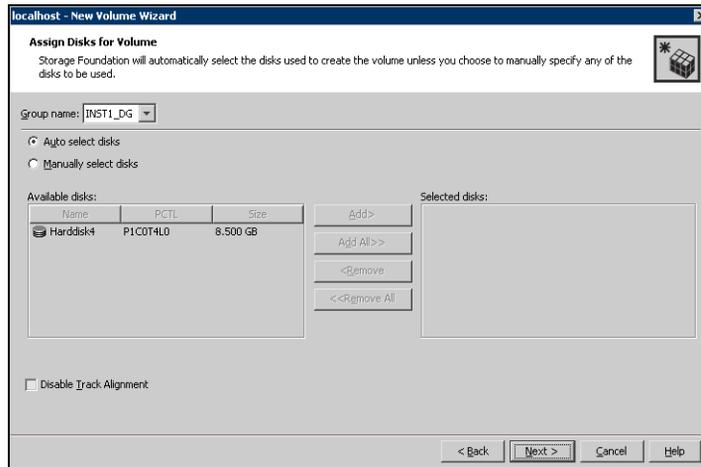
Warning: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

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

To create dynamic volumes

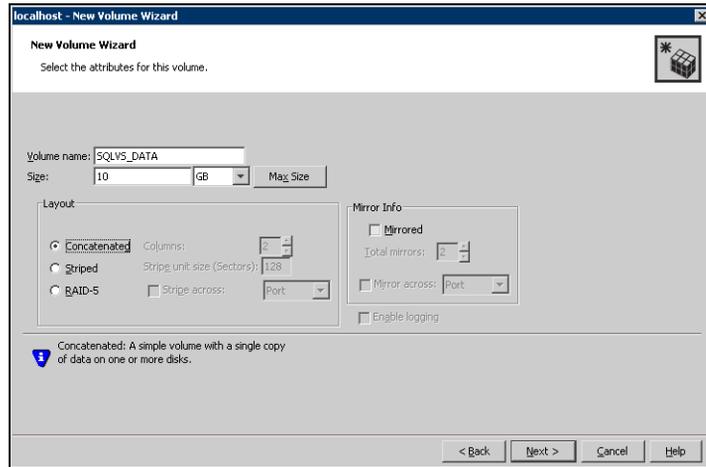
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

6 Select the disks for the volume.



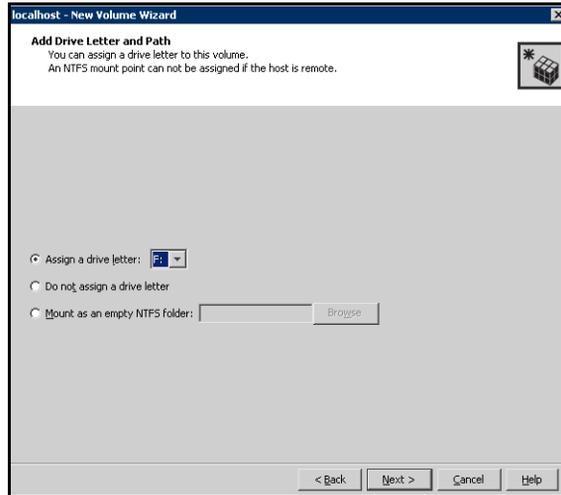
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



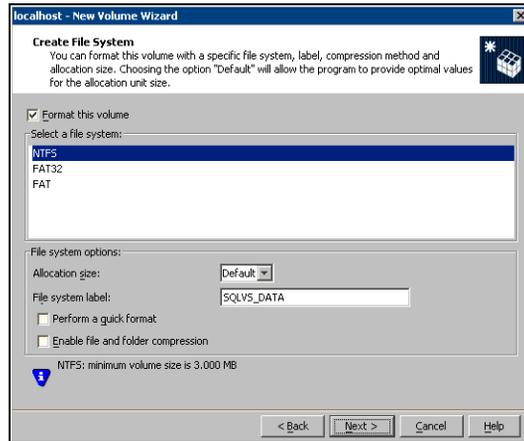
- Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - Provide a size for the volume.
 - If 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.
 - Verify that **Enable Logging** 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.
 - For the Replicator Log volume only:
Select **Do not assign a drive letter**.
- 9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
- For the Replicator Log volume only: Clear the Format this volume check box.
- 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**.

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

12 Repeat these steps to create additional volumes.

Installing and configuring SQL Server 2000 on the first node

Before installing Microsoft SQL Server 2000, ensure the cluster disk group is imported to the first node and the volumes are mounted.

Complete the following procedures to install and configure Microsoft SQL Server 2000:

- [Installing Microsoft SQL Server](#)
- [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 591 and “[Adding drive letters to mount the volumes](#)” on page 591.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.

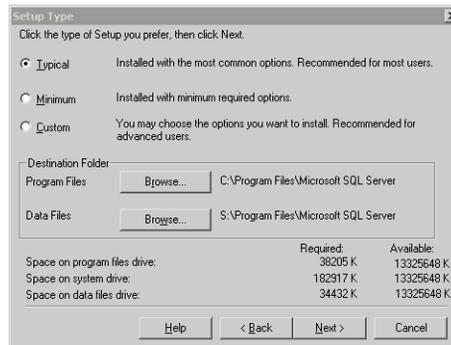
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

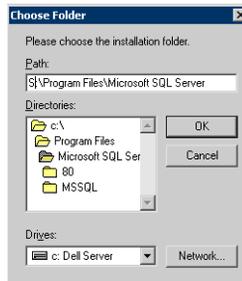
To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.

- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

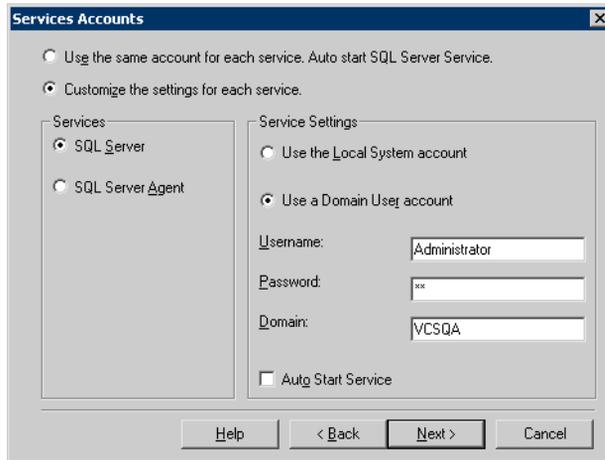


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 service”](#) on page 590
- [“Deporting the cluster disk group”](#) on page 590
- [“Importing the cluster disk group”](#) on page 591
- [“Adding drive letters to mount the volumes”](#) on page 591
- [“Renaming shared SQL Server 2000 files”](#) on page 593

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.

- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) on the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.

In the **Import Dynamic Disk Group** dialog box, click **OK**.

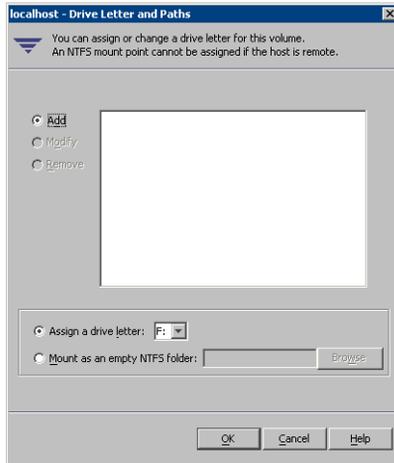
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 593
- [“Removing shared SQL Server files”](#) on page 596

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 591 and [“Adding drive letters to mount the volumes”](#) on page 591.

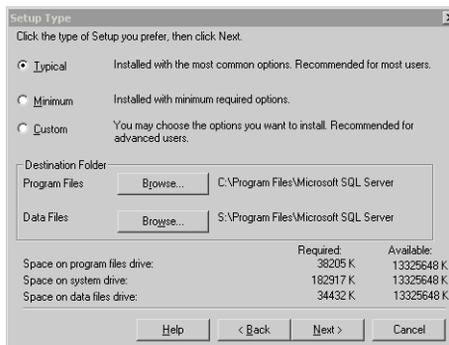
Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

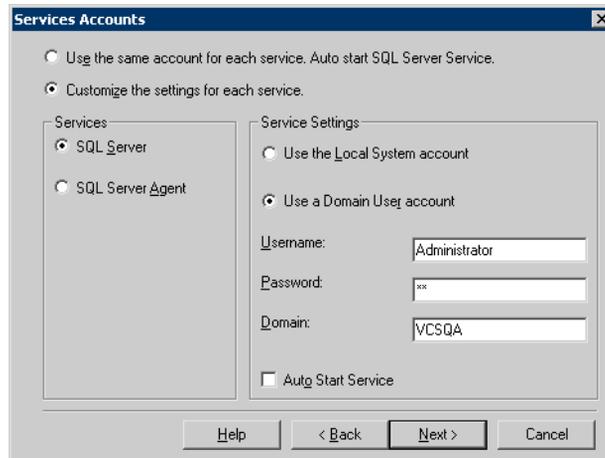
- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.

- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node](#)” on page 589 and “[Installing SQL Server 2000 on the second node](#)” on page 593 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

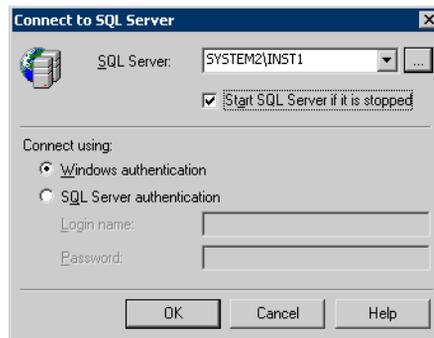
Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

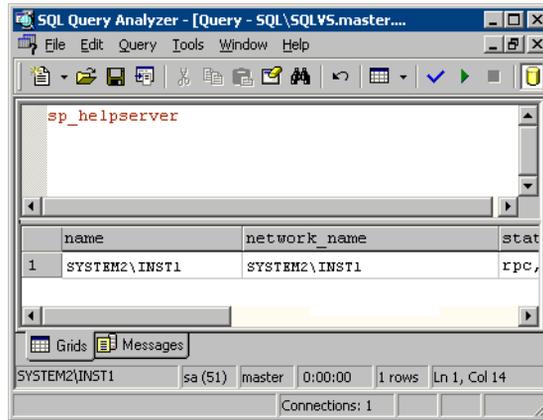
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example *SYSTEM2\INST1*.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
“sp_dropserver *\System_Name\Instance_Name*.’”
where *System_Name\Instance_Name* is the name noted in step 3.
For example, for named instance:
“sp_dropserver \SYSTEM2\INST1.’”
For example, for a default instance:
“sp_dropserver \SYSTEM1.’”
 - Press F5.
- 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter
`"sp_addserver 'Virtual_Server_Name\Instance_Name', local"`
For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.
 - Press F5.

Configuring the VCS SQL Server service group

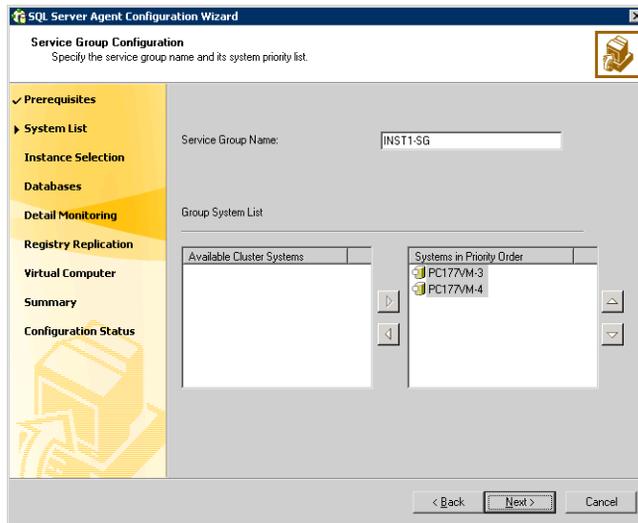
The SQL Server Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Complete the following tasks before configuring the service group:

- Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See [“Installing Veritas Storage Foundation HA for Windows”](#) on page 552.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 560.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- Ensure that you are a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify that the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. You specify this name when configuring the service group.
- Assign a unique virtual IP address to the SQL Server 2000 instance. You specify this IP address when configuring the service group.

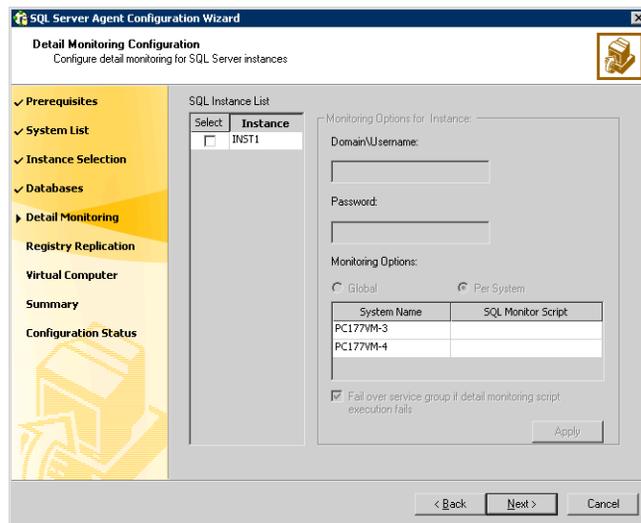
To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

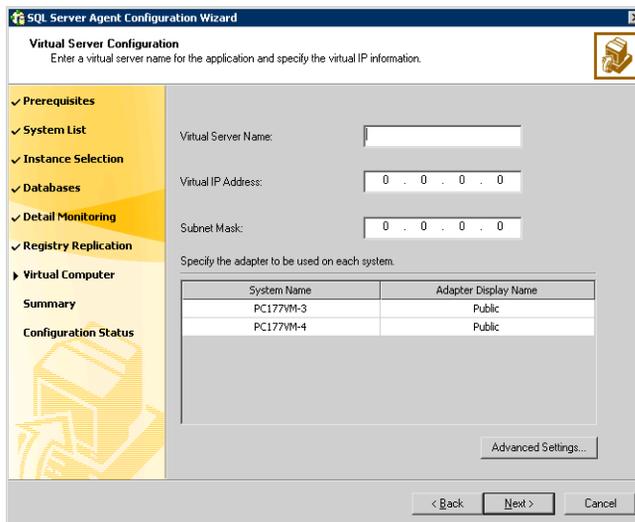
- Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
 Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 601 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to

fail over the service group if the detail monitoring script execution fails.

- Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 601 for all the instances for which detail monitoring will be configured.
- 10 Click **Next**.
- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

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

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
 - 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
 - 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
 - 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

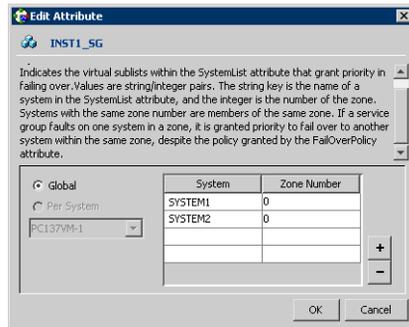
The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the SQL Server service group (INST1_SG) in the left pane and the Properties tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.



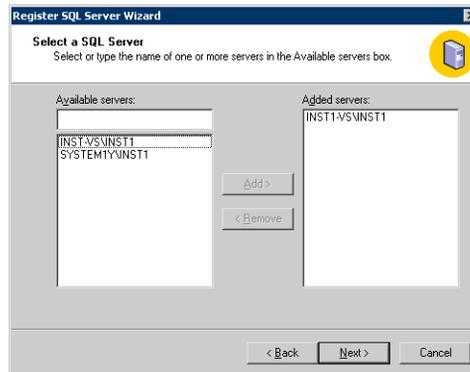
- 7 Click **OK**.

Registering the virtual server in the SQL Server

In the SQL Server Enterprise Manager, register the VCS virtual server (INST1-VS) and delete the physical server name.

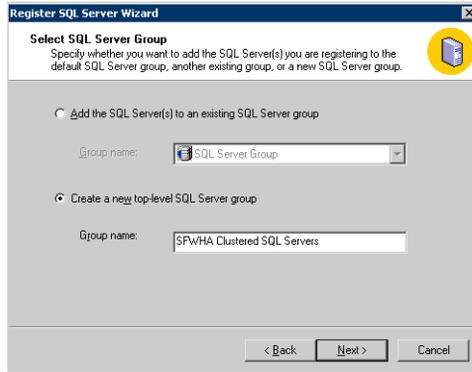
To register the VCS virtual server

- 1 Click **Start > All Programs > Microsoft SQL Server > Enterprise Manager**.
- 2 In the right pane, expand the Microsoft SQL Servers.
- 3 Right-click the **SQL Server Group** and select **New SQL Server Registration**.
- 4 On the Welcome screen, click **Next**.
- 5 In the **Select a SQL Server** window, select the virtual server\SQL instance (for example INST1-VS\INST1).



- 6 Click the **Add** button.
- 7 Click **Next**.
- 8 In the **Select an Authentication Mode** window, select an authentication mode and click **Next**.

- 9 In the **Select SQL Server Group** window, select Create a new top-level SQL Server group, or leave it in the existing group. If a new SQL Server group is created, enter the **Group name**.



- 10 Click **Next**.
- 11 In the **Completing the Register the SQL Server Wizard** window, click **Finish**.
- 12 Wait for the **Register SQL Server Messages** window that shows the server registration completed successfully and click **Close**.

To delete the physical server name

- 1 In the Microsoft SQL Server **Enterprise Manager**, expand the Microsoft SQL Servers and the **SQL Server Group**.
- 2 Right-click the physical name of the SQL Server and instance (for example SYSTEM1\INST1) and select **Delete SQL Server Registration**.
- 3 Click **Yes** in the warning message window.
- 4 Verify the SQL servers in the Microsoft SQL Server **Enterprise Manager** window. In the left pane, the **Microsoft SQL Servers** should show the virtual server under **SFWHA Clustered SQL Servers** and no items under the **SQL Server Group**.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 580 for information on how to use the VEA console to create a volume.

Note: Best practice is to create a separate disk group with separate volumes for the user-defined database files.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.

- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as **CRITICAL**.

- Click **Next** to configure another SQL service group or an MSDTC service group.

Verifying the installation in the primary zone

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

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

To shut down an active cluster node

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

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

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, use the guidelines in the following list to complete the same tasks in the secondary zone.

Before you begin to configure the secondary zone, offline the following resources in the SQL service group in the primary zone:

- SQL Server resource (<sqlservicegroupname> - SQLServer2000)
- SQL Virtual Server name resource (<sqlservicegroupname> - Lanman)
- SQL Virtual IP resource (<sqlservicegroupname> - IP)

The remaining resources should be online, including the VMDg resources and the MountV resources.

- [“Reviewing the prerequisites”](#) on page 543
- [“Reviewing the configuration”](#) on page 547
- [“Configuring the storage hardware and network”](#) on page 549
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 552
- [“Configuring VxSAS”](#) on page 558
- [“Adding the systems in the secondary zone to the cluster”](#) on page 611
- [“Configuring cluster disk groups and volumes”](#) on page 577

During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:

- Cluster disk group name
- Volume sizes
- Volume names
- Drive letters
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 586
Select the same options in the secondary zone as you did for the primary zone.
- [“Preparing to install SQL Server on the second node”](#) on page 589
The instance name must be the same in the primary zone and secondary zone.

- [“Installing SQL Server 2000 on the second node”](#) on page 593

Note: After you install SQL Server 2000 on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes, to avoid conflicts during the configuration of the zones.

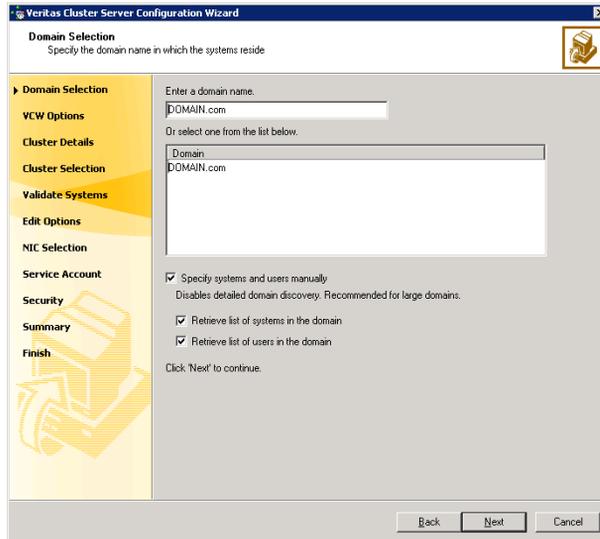
Adding the systems in the secondary zone to the cluster

Add the nodes in the secondary zone to the existing cluster with the following procedure.

To add a node to a VCS cluster

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

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



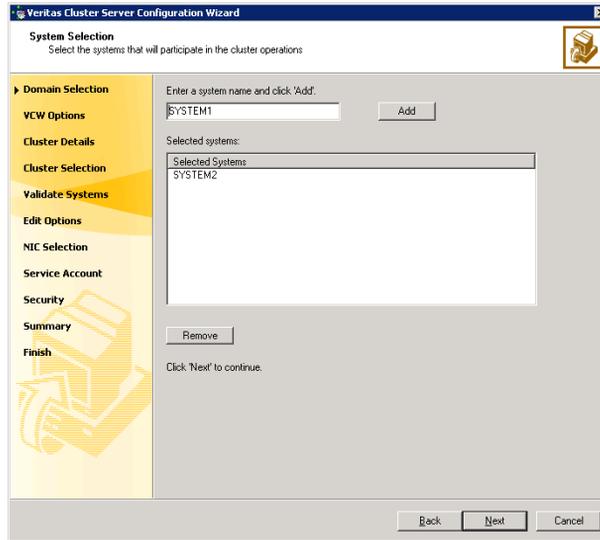
To discover information about all the systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 8](#) on page 615.

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

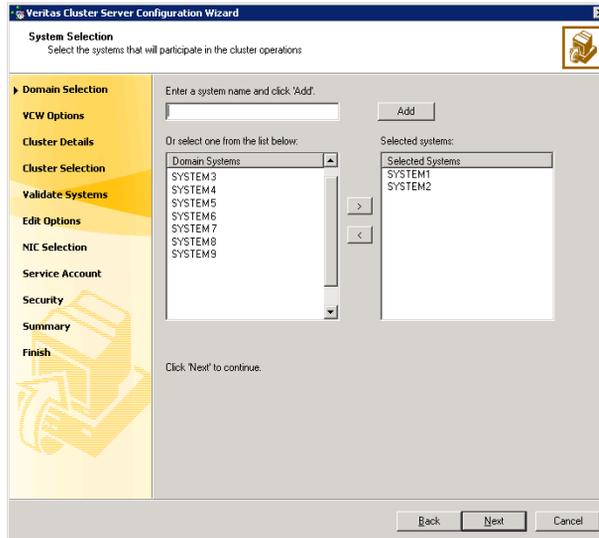
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 614. Otherwise proceed to the next step.

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



- Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
- If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 8](#) on page 615.

- 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.

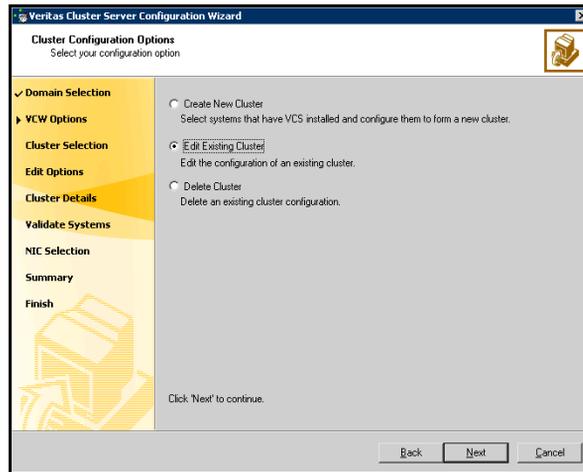
A system can be rejected for any of the following reasons:

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

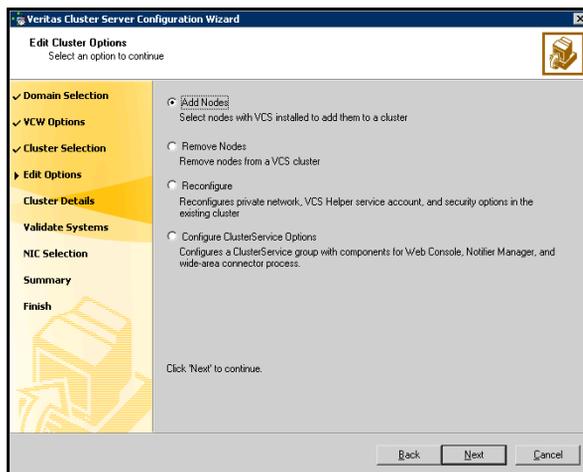
Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

Click **Next** to proceed.

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



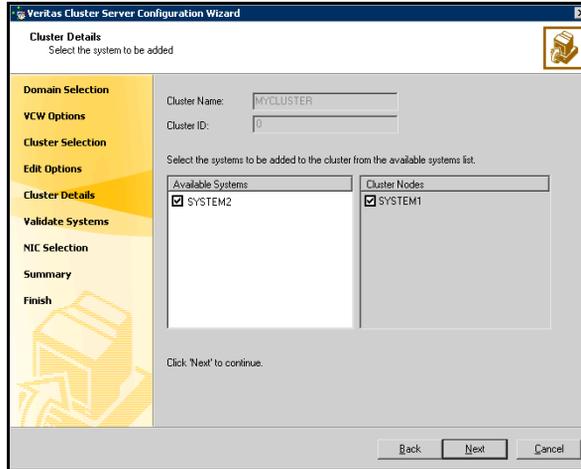
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#) on page 612, only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **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.

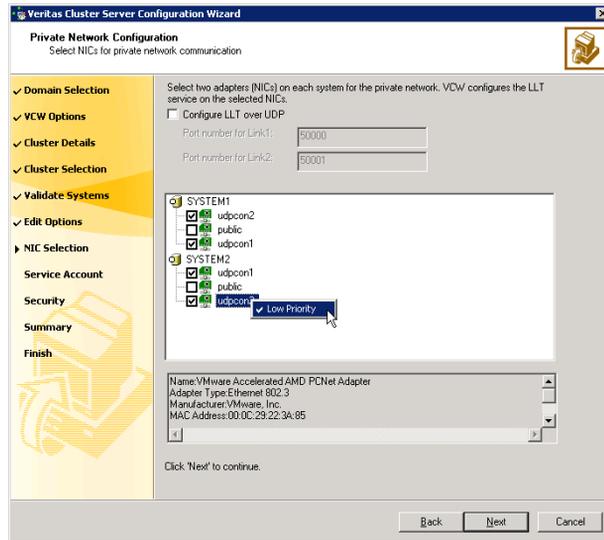
- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

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:

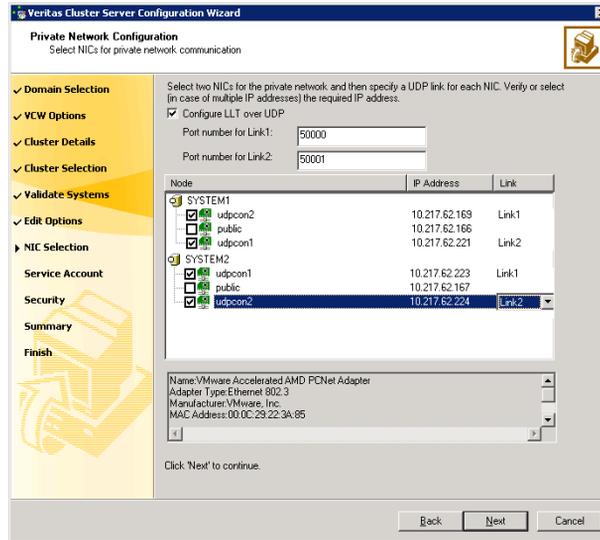
- To configure the VCS private network over Ethernet



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

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

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

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

- 15 Specify the 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**.

Return to the task list “[Creating a parallel environment in the secondary zone](#)” on page 610.

Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- Verify that the cluster disk group is imported and the volumes are mounted in the primary and secondary zone

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**. From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.

- 3 Read the Welcome page and click **Next**.
- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

Setup Replicated Data Set Wizard

Enter names for Replicated Data Set and Replicated Volume Group

Select the desired Primary host from the list of connected hosts.

Replicated Data Set name : INST1_RDS

Replicated Volume Group name : INST1_RVG

Primary Host : localhost

Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host.

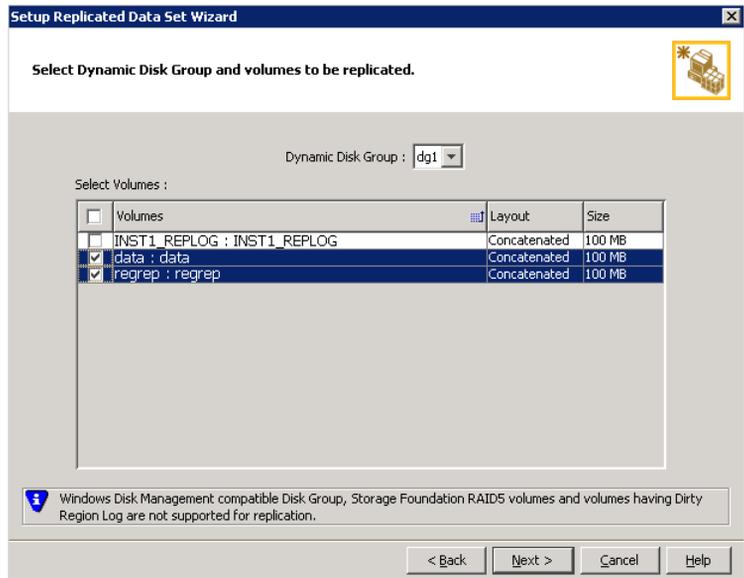
< Back Next > Cancel Help

By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

- 5 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

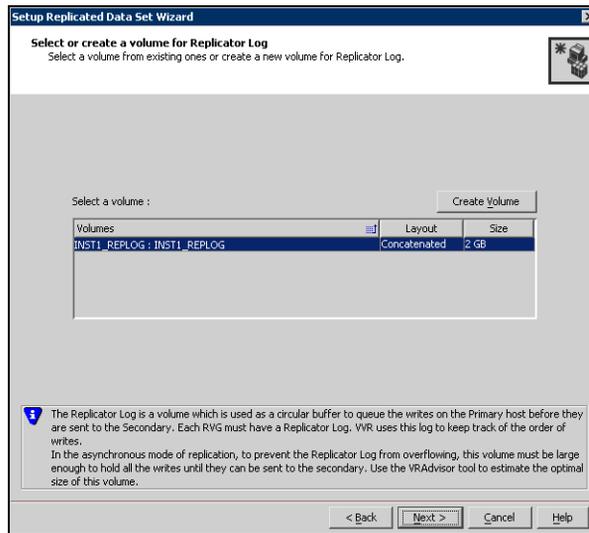


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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.

- 7 Click **Next**.

8 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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.

- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name Enter the name for the volume in the **Name** field.

Size Enter a size for the volume in the **Size** field.

Layout Select the desired volume layout.

- Disk Selection**
- Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.
 - Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select **Add** to move the disks into the selected disks pane.

- Click **OK** to create the Replicator Log volume.

- 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.

14 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows a window titled "Setup Replicated Data Set Wizard" with a sub-header "Edit replication settings". Below the sub-header is the instruction "Edit replication settings or click next." and a small icon of a server rack. The main area contains several configuration fields:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

Below these fields is an "Advanced" button. At the bottom left, there is a warning icon and the text "DHCP addresses are not supported by VWR." At the bottom right, there are four buttons: "< Back", "Next >", "Cancel", and "Help".

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not

wish to modify basic properties then replication can be started with the default values when you click **Next**.

- | | |
|---------------------------|--|
| Primary side IP | Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address. |
| Secondary side IP | Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address. |
| Replication Mode | <p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.</p> |
| Replicator Log Protection | The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows. |

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

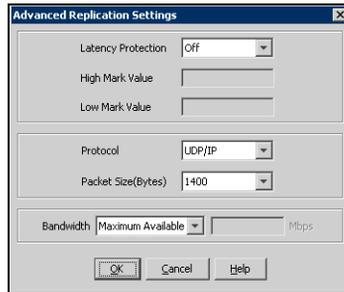
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

- **Off** is the default option and disables latency protection.
- **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
- **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
 Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.

18 Review the information.

Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

If the SQL Server user-defined database and files are in a separate disk group from the SQL Server system files, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 619 for the disk group (INST1_DB1_DG) that contains the SQL Server user-defined database and files. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name. See “[Sample configuration](#)” on page 548 for a list of example names.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication.

The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For additional information about service group types, see the *Veritas Cluster Server Administrator's Guide*.

Configure the RVG service group's resources manually by copying and modifying components of the SQL Server service group. Then create new RVG resources and bring them online.

The RVG service group for RDC contains the following resources:

Table 13-5 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the first disk group	Volume Manager disk group with SQL system files
VvrRvg for the first disk group	Replicated volume group with SQL system files
VMDg for the second disk group	Volume Manager disk group with SQL user-defined files
VvrRvg for the second disk group	Replicated volume group with SQL user-defined files

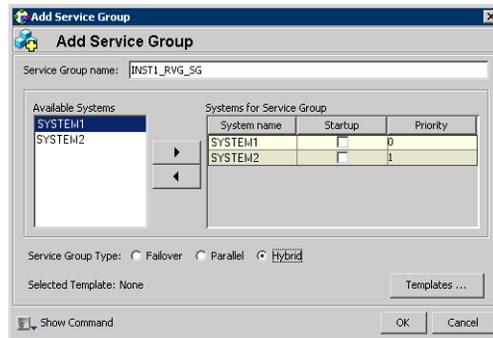
Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.

3 In the **Add Service Group** window:



- a Enter a name for the service group, for example INST1_RVG_SG.
- b Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- c Select **Hybrid**.
- d Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group’s resources manually for RVG by completing the following tasks:

- [“Configuring the IP and NIC resources”](#) on page 632
Copy IP and NIC resources of the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Configuring the VMDg resources for the first disk group”](#) on page 633 and [“Configuring the VMDg resources for the second disk group”](#) on page 635
Copy the VMDg resources for all disk groups in the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Adding the VVR RVG resources for the disk groups”](#) on page 636
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- [“Linking the VVR RVG resources to establish dependencies”](#) on page 639
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk groups. Configure the RVG service group’s VMDg resources to point to the disk groups that contain the RVGs.

- “Deleting the VMDg resource from the SQL Server service group” on page 641
Delete the VMDg resources from the SQL Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources

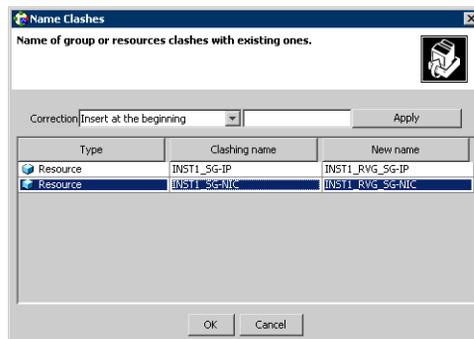
Configure the following resources and attributes for the IP and NIC:

Table 13-6 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the IP resource (INST1_SG-IP), and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 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 (INST1_RVG_SG-IP) 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 VVR IP address for the Primary Zone as the scalar value.
- 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 (INST1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (INST1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the first disk group

Create the VMDg resource in the SQL Server service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the SQL Server service group for the MountV resource:

Table 13-7 MountV resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
MountV (for the SQL Server system volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name
Resources for the disk group for the SQL user-defined database files:	
MountV (for the SQL Server user-defined database log)	VMDg Resource Name Volume Name
MountV (for the SQL Server user-defined database)	VMDg Resource Name Volume Name

To create the VMDg resource for the first disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the first disk group, with the SQL system files (INST1_SG-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server system data files (INST1_SG-MountV) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server system data files (INST1_DATA_FILES).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the registry volume (INST1_SG-MountV-1) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the registry volume (INST1_REGREP_VOL).

- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg) and select **Enabled**.

Configuring the VMDg resources for the second disk group

Repeat the VMDg and MountV configuration for the second disk group.

To create the VMDg resource for the second disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the second disk group, with SQL user-defined files (INST1_SG-VMDg-1), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg-1.
- 6 Click **OK**.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined log (INST1_SG-MountV-2) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined log (INST1_DB1_LOG).

- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg-1.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined database (INST1_SG-MountV-3) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined database (INST1_DB1_VOL).
- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg-1) and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add VVR RVG resources for replication of the disk groups.

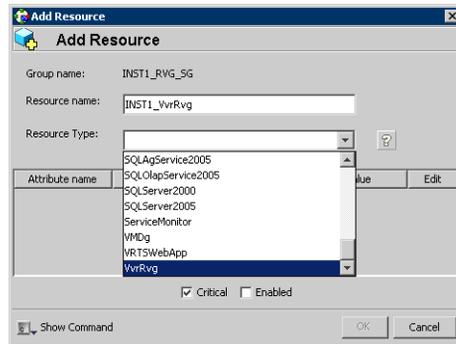
Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 13-8 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
VvrRvg	VMDgResName IPResName
Resources for the disk group for the SQL user-defined database files:	
VvrRvg	VMDgResName IPResName

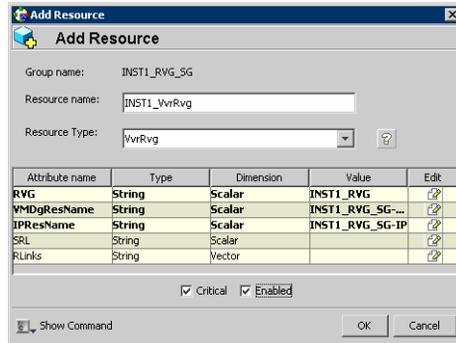
To create the VVR RVG resource for the first disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- Enter the **Resource Name** for the VVR RVG resource.
 - Select the **Resource Type** of VvrRvg.
- 3 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
 - 4 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_RVG.
 - 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, for example INST1_RVG_SG-VMDg.
 - 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, for example INST1_RVG_SG-IP.
 - 11 Click **OK**.

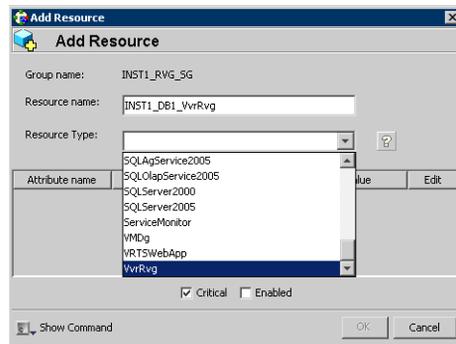
12 In the **Add Resource** window, verify that the attributes have been modified:



13 Click **OK**.

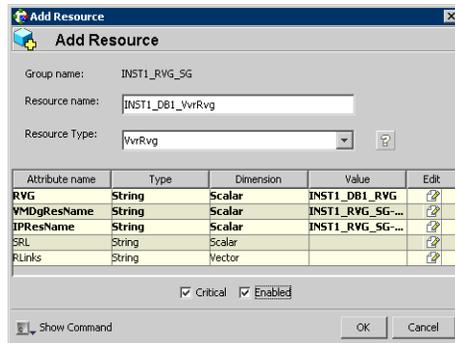
To create the VVR RVG resource for the second disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- Enter the **Resource Name** for the VVR RVG resource for the second disk group.
 - 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 group that is being managed, for example INST1_DB1_RVG.
 - 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 disk group containing the RVG, for example INST1_RVG_SG-VMDg-1.
- 8 Click **OK**.
- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name IP resource managing the IP address for replication, for example INST1_RVG_SG-IP. In this example both disk groups are using the same IP resource for replication.
- 11 Click **OK**.
- 12 In the **Add Resource** window, verify that the attributes have been modified:



- 13 Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the VVR RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:

Table 13-9 Dependencies for VVR RVG resources for RDC

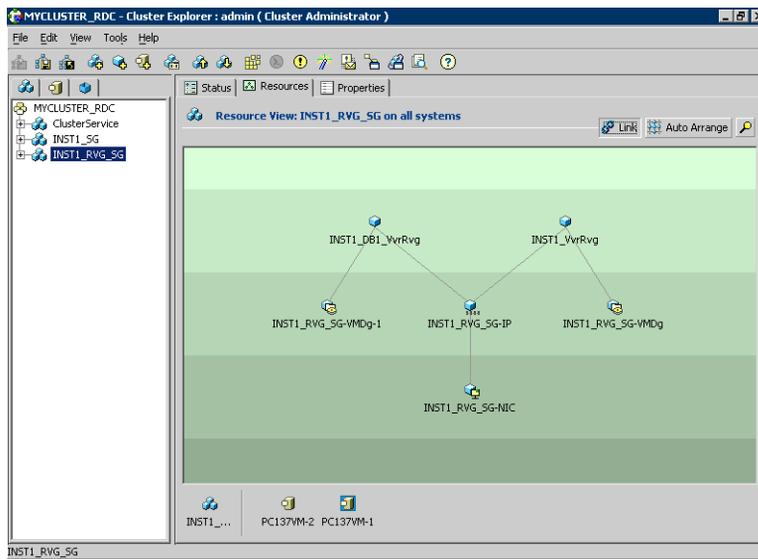
Parent	Child
Resources for the disk group for the SQL system files:	
INST1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_ VvrRvg	The VMDg for the SQL system files, for example INST1_RVG_SG-VMDg
Resources for the disk group for the SQL user-defined database files:	

Table 13-9 Dependencies for VVR RVG resources for RDC (Continued)

Parent	Child
INST1_DB1_VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_DB1_VvrRvg	The VMDg for the SQL user-defined database files, for example INST1_RVG_SG-VMDg-1

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example INST1_DB1_VvrRvg.
- 4 Click the child resource, for example INST1_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources:



Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg resource from the SQL Server service group

The VMDg resources must now be manually deleted from the SQL Server service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) from the left pane.
- 2 In the **Resources** tab display area, right-click the VMDg resource for the first disk group (INST1_SG-VMDg) and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the **Resources** tab display area, right-click the VMDg resource for the second disk group (INST1_SG-VMDg-1) and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVGPrimary to the SQL Server service group for each of the SQL Server disk groups (system and user-defined) and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the SQL Server service group for the RVG Primary resources:

Table 13-10 RVG Primary resources

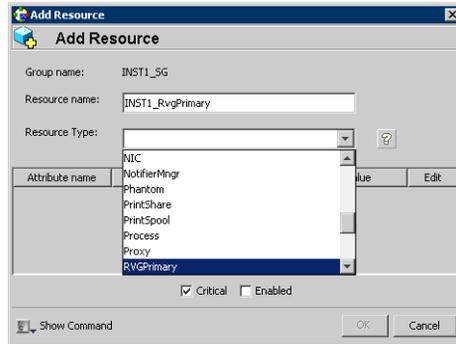
Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
RVGPrimary	RvgResourceName
Resources for the disk group for the SQL user-defined database files:	
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For all disk groups, create an RVG Primary Resource for replication.

To create the RVG Primary resource for the SQL Server system disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:



- 3 Enter the **Resource Name** for the RVG Primary resource for the SQL Server system files disk group, for example INST1_RvgPrimary.
- 4 Select the **Resource Type** of RVGPrimary.
- 5 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 6 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_VvrRvg and click **OK**.
- 7 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server 5.1 Administrator's Guide* for more information about the RVG Primary agent.
- 8 Verify that **Critical** and **Enabled** are both checked.
- 9 Click **OK**.

To create the RVG Primary resource for the SQL Server user-defined database disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:
 - Enter the **Resource Name** for the RVG Primary resource for the SQL Server user-defined database disk group, for example INST1_DB1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_DB1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.
- 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 SQL Server service group (INST1_SG) to establish the dependencies between the resources for replication.

Start from the top parent and link the following resources:

Table 13-11 Dependencies for the RVG Primary resources for RDC

Parent	Child
INST1_SG-MountV	INST1_RvgPrimary
INST1_SG-MountV-1	INST1_RvgPrimary
INST1_SG-MountV-2	INST1_DB1_RvgPrimary
INST1_SG-MountV-3	INST1_DB1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 Click the **Link** button in the right pane.

- 3 Click the parent resource, for example INST1_SG-MountV.
- 4 Click the child resource, for example INST1_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the SQL Server service group (INST1_SG) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource (INST1_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right click the second RVG Primary resource (INST1_DB1_RvgPrimary) and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG Service Group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.

Setting a dependency between the service groups

To ensure that the SQL Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the SQL Server service group.

The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

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 SQL Server service group (the parent service group), for example INST1_SG.
- 5 Click the RVG service group (the child resource), for example INST1_RVG_SG.
- 6 In the **Link Service Groups** window:



- Select the **Relationship** of **online local**.
- Select the **Dependency Type** of **hard**.
- Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

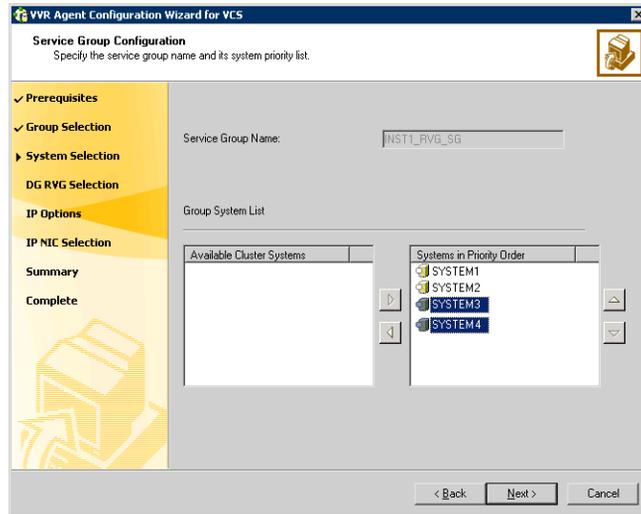
Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

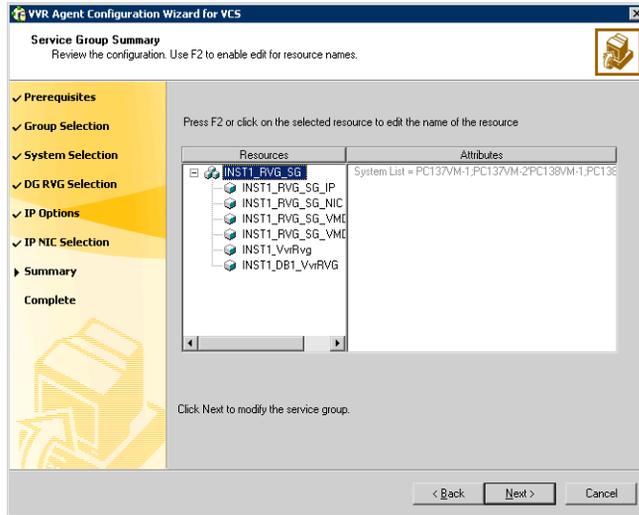
- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - 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:



- In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group’s system list.
 To remove a node from the service group’s system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
 - Click **Next**.
- 6 If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
 - 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
 - 8 In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
 - 9 If a VCS error appears, click **OK**.
 - 10 In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.

11 Review the summary of the service group configuration:



The **Resources** box lists the configured resources. 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**.

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 (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the IP resources for failover

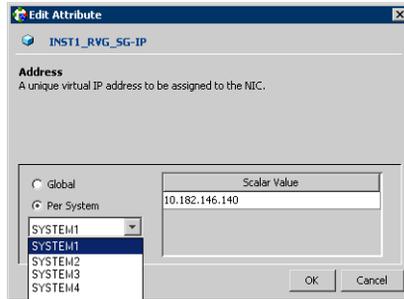
Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.

In the event of a system or SQL 2000 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group

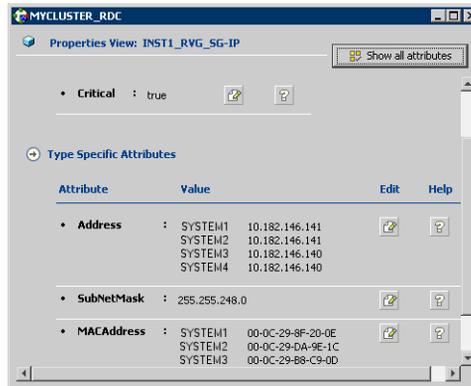
- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.

- 4 In the Edit Attributes window, edit the Address attribute.



- 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 (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone

have the same IP address, that is different from the IP address for the primary zone.



- 6 Close the Properties View window.
- 7 Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding the nodes from the secondary zone to the SQL Server service group

Use the SQL Server Agent Configuration Wizard to add the nodes from the secondary zone to the SQL Server service group.

To add the nodes from the secondary zone to the SQL Server service group

- 1 Select **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Configuration Wizard**.
- 2 In the Select Configuration Option dialog box, select **MS SQL Server - Service Group Configuration**. Select **Edit**, and click **Next**.
- 3 Verify that you have met the prerequisites listed and click **Next**.
- 4 Select the service group to be modified (INST1_SG) and click **Next**.
- 5 If a VCS notice message appears indicating that resources are online, click **Yes** to continue.
- 6 On the Service Group Configuration page, select the nodes in the secondary zone, use the arrow button to move them from **Available Cluster Systems** to **System in Priority Order**.

To change the priority of a system in the **Systems in Priority Order** list, select the system and click the up and down arrow icons. Arrange the

systems in priority order in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online.

This set of nodes selected for the SQL Server 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.

- 7 Click **Next**.
- 8 On the SQL Server Instance Selection page, click **Next**.
- 9 The User Databases List page summarizes the databases for this instance of SQL. Click **Next**.
- 10 On the Detail Monitoring Configuration page, clear the box in the SQL Instance List to disable monitoring, as required. Detailed monitoring is not necessary. Click **Next**.
- 11 On the Registry Replication Path page, click **Next**.
- 12 On the Virtual Server Configuration page, verify that the public adapter is used on each system and click **Next**.
- 13 In the Service Group Summary, review the service group configuration and click **Next**.
- 14 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.
- 15 Click **Finish**.

Configuring the zones in the SQL Server service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the SQL Server service group

- 1 From VCS Cluster Explorer, in the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.

- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- [Bringing the service group online](#)
- [Switching online nodes](#)

Bringing the service group online

After completing all configuration, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the Exchange Server service group online in the primary zone.

To bring the Exchange service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange Server service group (EVS1_SG1).
- 2 Click **Online**.

Switching online nodes

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

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) (**Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**).
- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Additional instructions for GCO disaster recovery

After completing the tasks for setting up a replicated data cluster for SQL Server 2000, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, if a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment. The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up VVR replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See “[Deploying disaster recovery: New SQL Server 2000 installation](#)” on page 779.

Configuring Replicated Data Clusters for SQL 2005

This chapter includes the following topics:

- [“Tasks for configuring Replicated Data Clusters for SQL 2005”](#) on page 656
- [“Reviewing the prerequisites”](#) on page 659
- [“Reviewing the configuration”](#) on page 664
- [“Configuring the storage hardware and network”](#) on page 666
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 668
- [“Configuring the cluster”](#) on page 678
- [“Configuring cluster disk groups and volumes”](#) on page 694
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 703
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 708
- [“Installing SQL Server 2005 on the second node”](#) on page 711
- [“Setting the internal name of the clustered instance”](#) on page 715
- [“Configuring the VCS SQL Server service group”](#) on page 718
- [“Creating the primary system zone”](#) on page 723
- [“Creating a SQL Server user-defined database”](#) on page 724
- [“Verifying the installation in the primary zone”](#) on page 726
- [“Creating a parallel environment in the secondary zone”](#) on page 727
- [“Adding the systems in the secondary zone to the cluster”](#) on page 728

- [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 736
- [“Configuring a hybrid RVG service group for replication”](#) on page 747
- [“Setting a dependency between the service groups”](#) on page 762
- [“Adding the nodes from the secondary zone to the RDC”](#) on page 763
- [“Verifying the RDC configuration”](#) on page 770
- [“Additional instructions for GCO disaster recovery”](#) on page 771

Tasks for configuring Replicated Data Clusters for SQL 2005

Configure the high availability and SQL 2005 components on the primary and secondary zones. Then complete the Replicated Data Set solution by configuring the components for both zones.

Refer to the *Veritas Volume Replicator Administrator’s Guide* for additional details on VVR.

The table below outlines the high-level objectives and the tasks to complete each objective.

Table 14-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Reviewing the prerequisites” on page 659	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 664	<ul style="list-style-type: none">■ Understanding active-passive configuration and zone failover in a RDC environment■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 666	<ul style="list-style-type: none">■ Setting up the storage hardware for a cluster environment■ Verifying the DNS entries for the systems on which SQL will be installed

Table 14-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 668	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ■ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ■ Configuring VxSAS
“Configuring the cluster” on page 678	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 694	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2005 on the first node” on page 703	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2005 ■ Setting SQL Server services to manual start ■ Configuring SQL services
“Preparing to install SQL Server 2005 on the second node” on page 708	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters ■ Removing shared SQL files from the cluster disk group

Table 14-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Installing SQL Server 2005 on the second node” on page 711	Installing and configuring SQL Server 2005
“Setting the internal name of the clustered instance” on page 715	Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 718	Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating the primary system zone” on page 723	<ul style="list-style-type: none"> ■ Creating the primary system zone ■ Adding the nodes to the primary zone
“Creating a SQL Server user-defined database” on page 724	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server ■ Adding resources for a user-defined database in VCS
“Verifying the installation in the primary zone” on page 726	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“Creating a parallel environment in the secondary zone” on page 727	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Configuring disk groups and volumes for SQL
“Setting up the Replicated Data Sets (RDS)” on page 736	Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones
“Configuring a hybrid RVG service group for replication” on page 747	<ul style="list-style-type: none"> ■ Creating a hybrid Replicated Volume Group (RVG) service group ■ Configuring the hybrid RVG service group

Table 14-1 Tasks for configuring Replicated Data Clusters for SQL 2005

Objective	Tasks
“Setting a dependency between the service groups” on page 762	Setting up a dependency from the VVR RVG Service Group to the SQL Server Service Group
“Adding the nodes from the secondary zone to the RDC” on page 763	<ul style="list-style-type: none"> ■ Adding the nodes from the secondary zone to the RVG service group ■ Configuring the IP resources for failover ■ Adding the nodes from the secondary zone to the SQL Server service group
“Verifying the RDC configuration” on page 770	Verifying that failover occurs first within zones and then from the primary to the secondary zone

Reviewing the prerequisites

Verify that the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation. This replication recovery solution requires installation and configuration at a primary zone and a secondary zone.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at:

<http://www.symantec.com/business/support/index.jsp>

Supported software

The following software is supported:

- Veritas Storage Foundation HA 5.1 for Windows (SFW HA)
 Select the installation options for the Veritas Cluster Server Database Agent for Microsoft SQL and the Veritas Volume Replicator.
 For a Disaster Recovery configuration, also select the Global Clustering Option.

- For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | ■ Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) |
| | ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) |
| | ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| | ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) |
| | ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| | ■ Windows Server 2008 for 64-bit Itanium (IA64) |
| | ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | ■ Windows Server 2008 for 64-bit Itanium (IA64) |
| | ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- Shared disks to support applications that migrate between nodes in the cluster. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs. One shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 663.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL Virtual Server.
 - A minimum of one static IP address for each physical node in the cluster.
 - One static IP address for each replication site.
 - For VVR, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on `C:\WINDOWS` of one node,

installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.
- When installing, install only in a single domain.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Note: Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `xclus UseSystemBus ON` command

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 14-2](#) on page 664 estimates disk space requirements for SFW HA.

Table 14-2 Disk space requirements

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

Reviewing the configuration

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the 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

You should have these IP addresses available before you start deploying your environment.

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:

Primary zone

SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
INST1_SG	Microsoft SQL Server 2005 service group
INST1-VS	Virtual SQL Server cluster

INST1	SQL Instance Name
INST1_DG	Cluster disk group for SQL system database and files
INST1_DATA_FILES	Volume for Microsoft SQL Server system data files
INST1_REGREP_VOL	Volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	Replicator log volume required by VVR
INST1_DB1_DG	Cluster disk group for SQL Server user-defined database and files
INST1_DB1_VOL	Volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	Volume for storing a Microsoft SQL Server user-defined database log file
INST1_DB1_REPLOG	Replicator log volume required by VVR for SQL user-defined database

Secondary zone

SYSTEM3 & SYSTEM4 First and second nodes of the secondary zone

All the other parameters are the same as on the primary zone.

**RDS and VVR
 Components**

INST1_RDS	RDS name for SQL system database and files
INST1_RVG	RVG name for SQL system database and files
INST1_RVG_SG	Replication service group for SQL system database and files
INST1_DB1_RDS	RDS name for SQL Server user-defined database and files
INST1_DB1_RVG	RVG name for SQL Server user-defined database and files
INST1_DB1_RVG_SG	Replication service group for SQL Server user-defined database and files

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
 When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the options to install VVR and the Veritas Cluster Server Database Agent for SQL. The Veritas Cluster Server Enterprise Agent for VVR is automatically installed with the VVR installation.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

Table 14-3 on page 669 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 14-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 670.

To change the driver signing options on each local system

- 1 Open the Control Panel and click **System**.
- 2 Click the Hardware tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or one of the other options from the table, to allow installation to proceed.
- 4 Click **OK**.
- 5 Repeat for each computer.

If you do not change these options, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing options to their previous states.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install on Windows 2008 will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

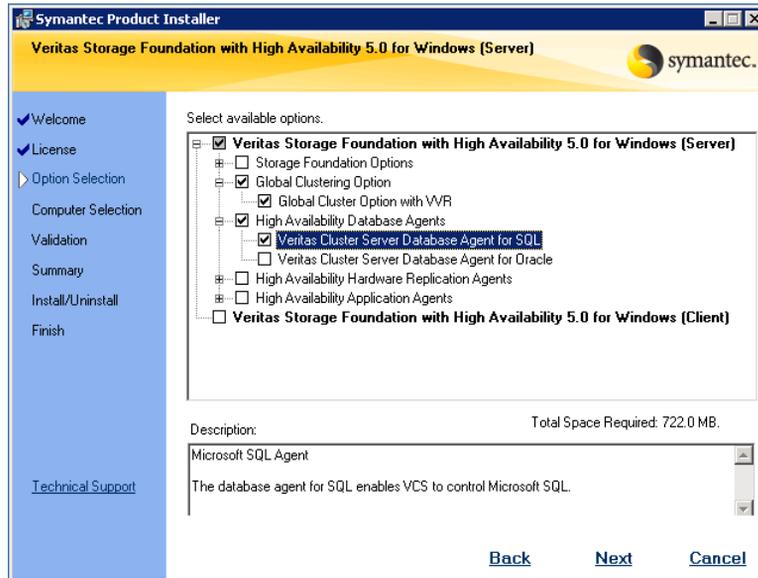
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**.
- 2 Choose the language for your installation and click **OK**. The SFW Select Product screen appears.
- 3 Click **Storage Foundation HA 5.1 for Windows**.

4 Do one of the following:



- Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 5 Review the Welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and click **Next**.
- 7 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
 If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 8 Repeat for additional license keys. Click **Next**
- To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 9 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:

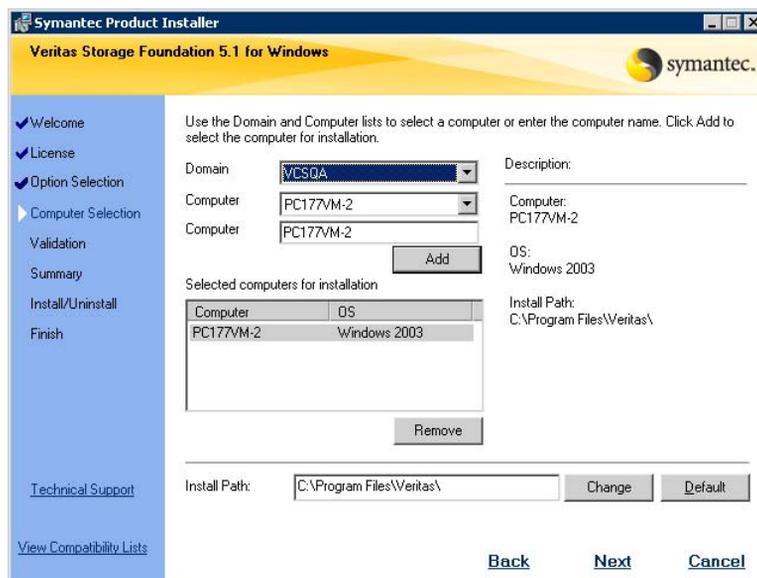


Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

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

Veritas Volume Replicator To use VVR for replication, you must select the option to install VVR.

10 Select the domain and the computers for the installation and click **Next**.



- Domain** Select a domain from the list.
- Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.
- Computer** To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.
- To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.
- Click a computer's name to see its description.
- Install Path** Optionally, change the installation path.
- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
 - To restore the default path, select a computer and click **Default**.
- The default path is:
 C:\Program Files\Veritas
 For 64-bit installations, the default path is:
 C:\Program Files (x86)\Veritas

- 11 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 12 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 13 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 14 Click **OK**.
- 15 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 16 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 17 When the installation completes, review the summary screen and click **Next**.
- 18 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 19 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 20 Review the log files and click **Finish**.
- 21 Click **Yes** to reboot the local node.

Configuring VxSAS

You can run the VVR Service Configuration (VxSAS) wizard after you install SFW HA on both the primary and secondary nodes. When you run the wizard, you can then specify the primary and secondary sites in one step.

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxsascfg.exe` from the command prompt of the required machine.

Read the information provided on the Welcome page and click **Next**.

- 2 Complete the Account Information panel as follows:

Account name Enter the administrative account name.
 (domain\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

Click **Next**.

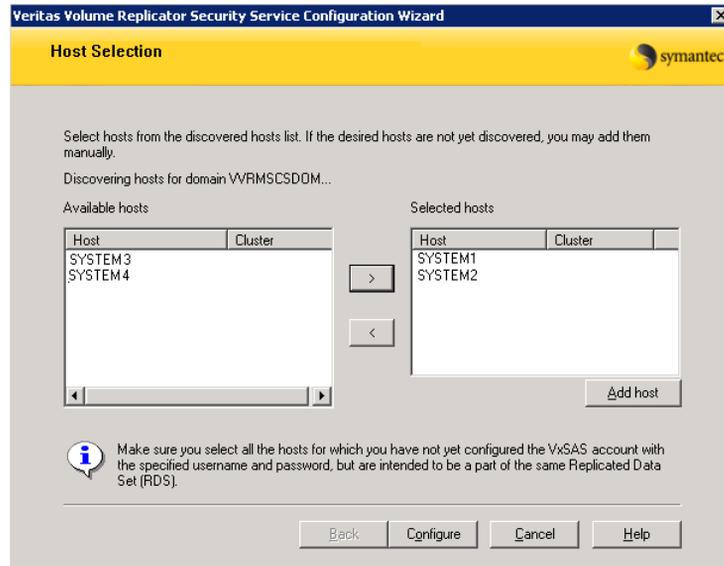
- 3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

- 4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

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

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

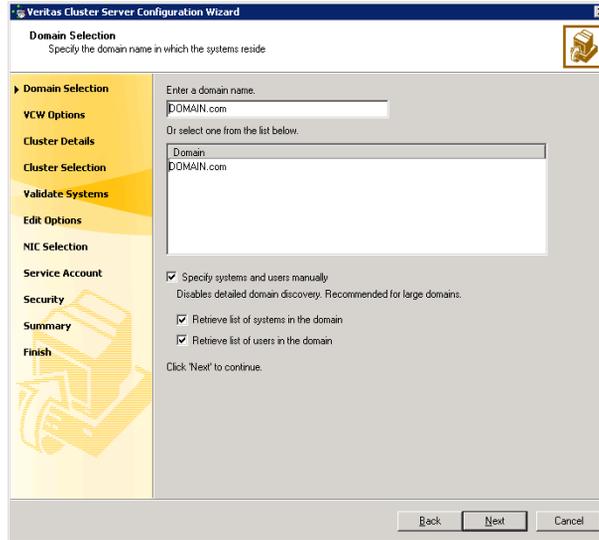
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Note: Add only systems in the primary zone (zone 0) to the cluster at this time.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



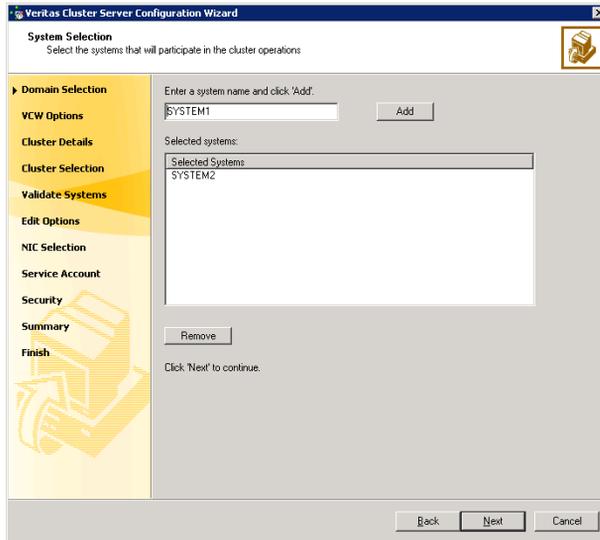
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
 Proceed to [step 8](#) on page 681.

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

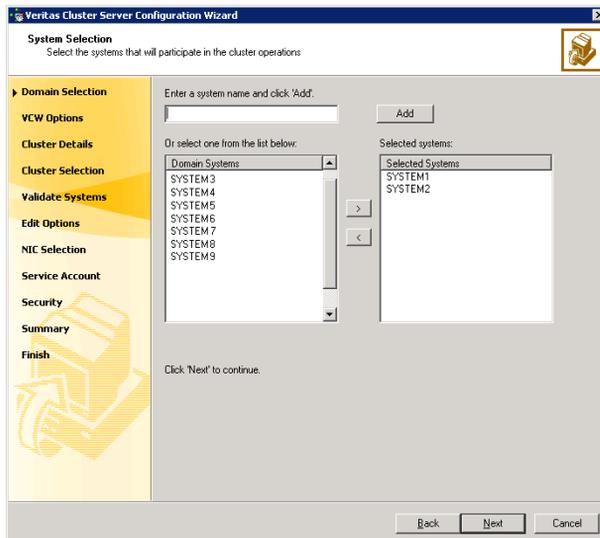
- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
 If you chose to retrieve the list of systems, proceed to [step 6](#) on page 680. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 681.

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



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

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

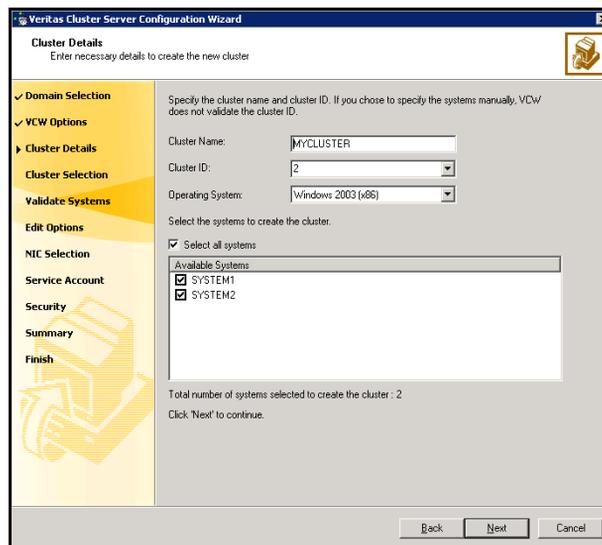
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 679 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

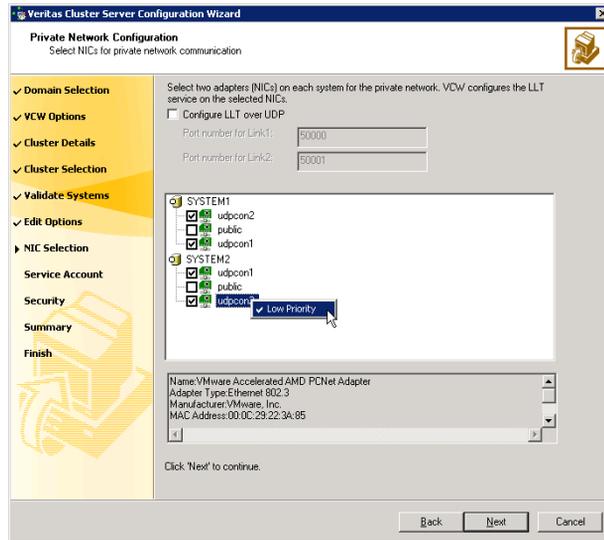
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 681, proceed to the next step. Otherwise, proceed to [step 12](#) on page 685.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

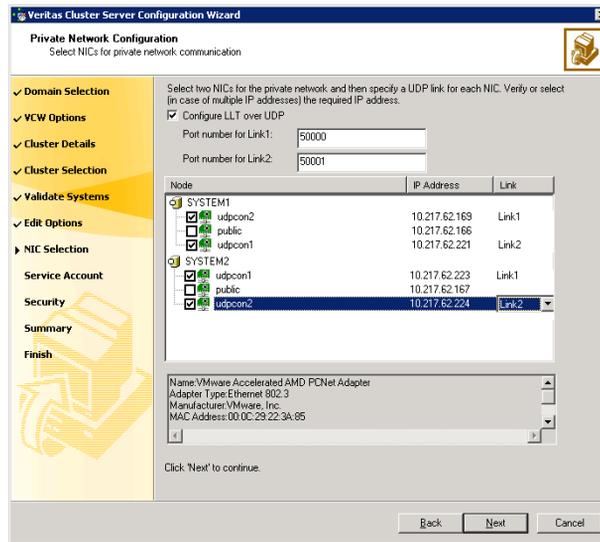
- To configure the VCS private network over Ethernet



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

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

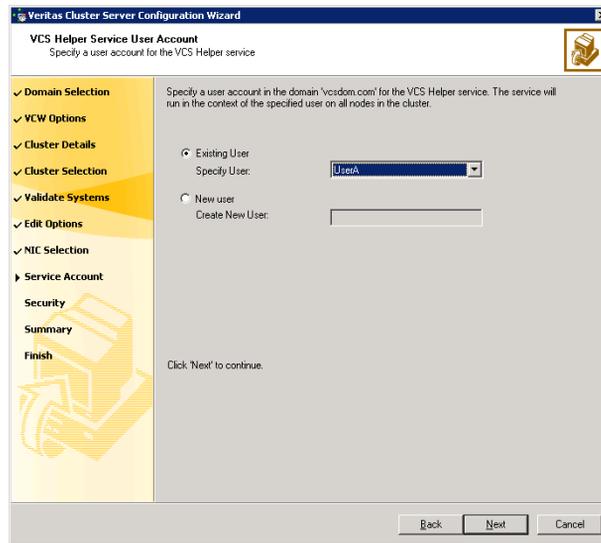
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



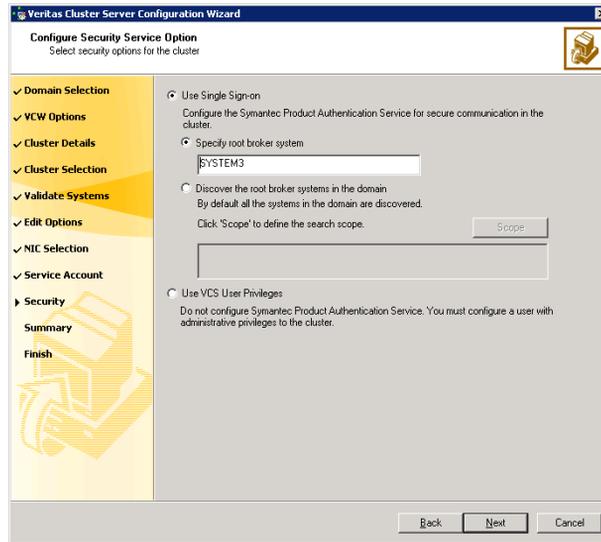
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 679, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
 Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

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

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 14-4](#) on page 687 contains some more examples of search criteria.

Table 14-4 Search criteria examples

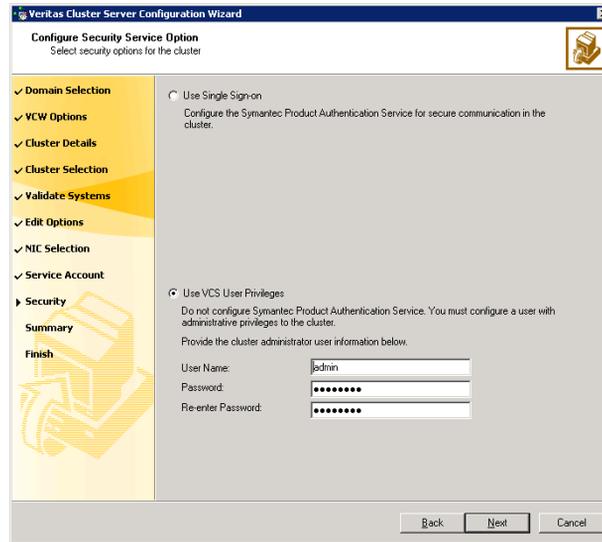
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.

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

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

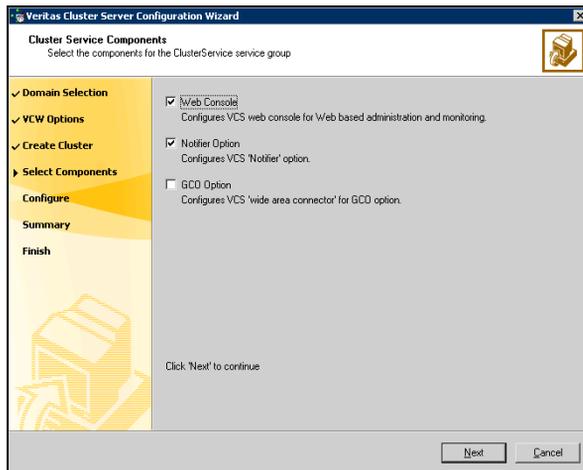
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 690.

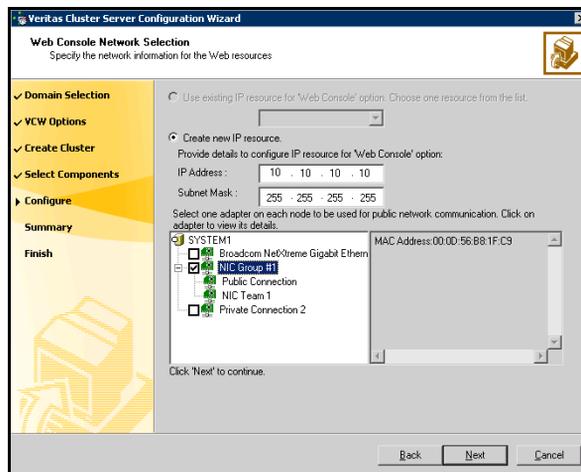
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See “[Configuring notification](#)” on page 691.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



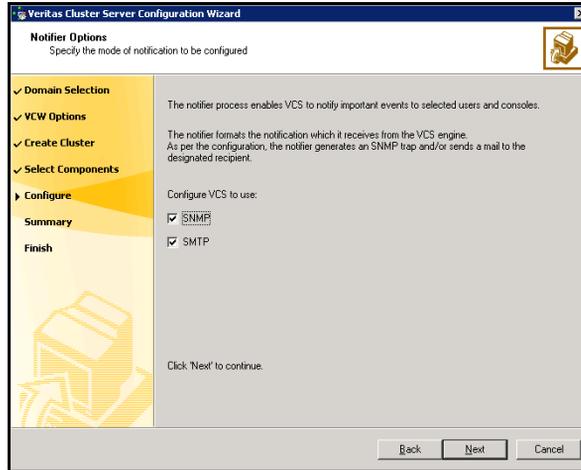
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 691.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

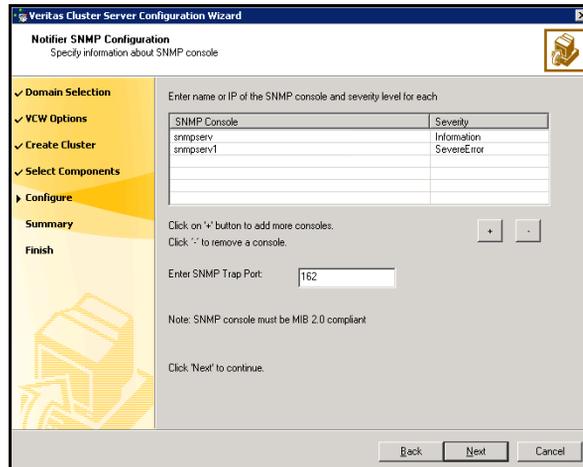
To configure notification

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

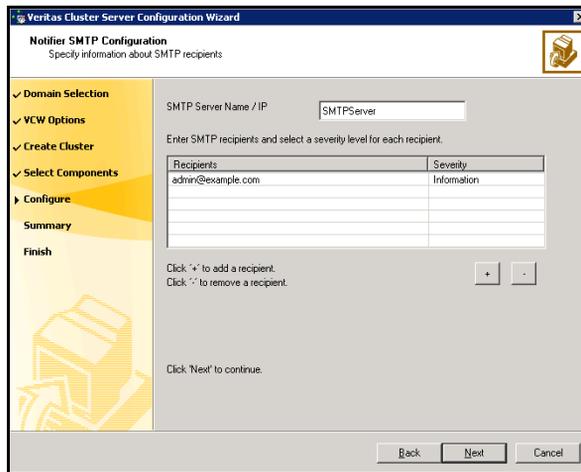


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

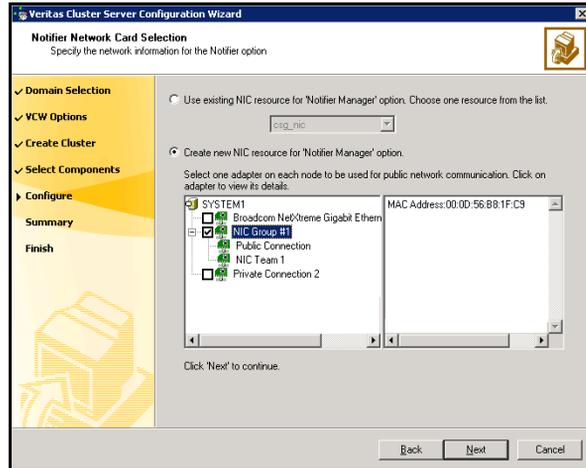


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

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

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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 disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the disk group are shared and are available from all nodes in each zone. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_REPLOG** contains the VVR Replicator Log. You can create this volume later while setting up the Replicated Data Sets. See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 736.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and files.

The following disk group and volumes may be created now or later in the configuration process.

- **INST1_DB1_DG** is the disk group for the SQL Server user-defined database and files
- **INST1_DB1_VOL** contains the user database files
- **INST1_DB1_LOG** contains the user database log files
- **INST1_DB1_REPLOG** contains the VVR Storage Replicator Log.

Warning: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Creating a cluster disk group

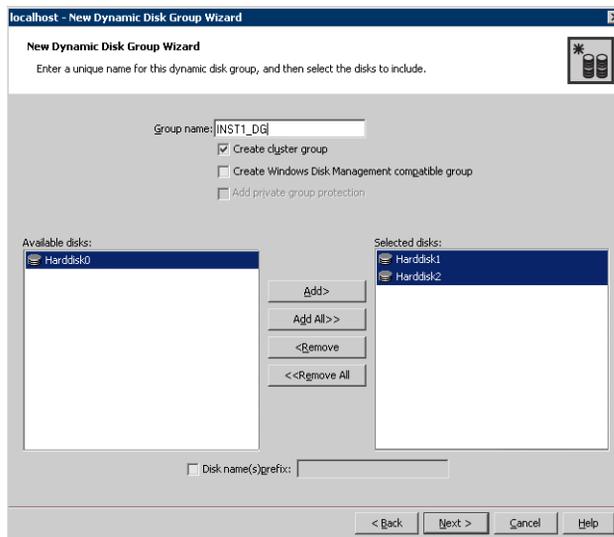
Create a cluster disk group on the first node of the cluster. You can repeat the process to create the disk group for the SQL Server user-defined database and files at this time or later.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes in the cluster disk group (INST1_DG) for the system files on the first node of the cluster:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- INST1_REPLOG contains the VVR Storage Replicator Log.
You can create the INST1_REPLOG volume at this time or during the process of creating Replicated Data Sets.
See “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 736.

You can create the INST1_DB1_VOL and INST1_DB1_LOG volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 724.

As a best practice, create a separate disk group and volumes for SQL Server user-defined database and the log files. The following volumes in the disk group (INST1_DB1_DG) for the SQL Server user-defined database and files may be created now or later in the configuration process:

- INST1_DB1_VOL contains the user database files
- INST1_DB1_LOG contains the user database log files
- INST1_DB1_REPLOG contains the VVR Storage Replicator Log.

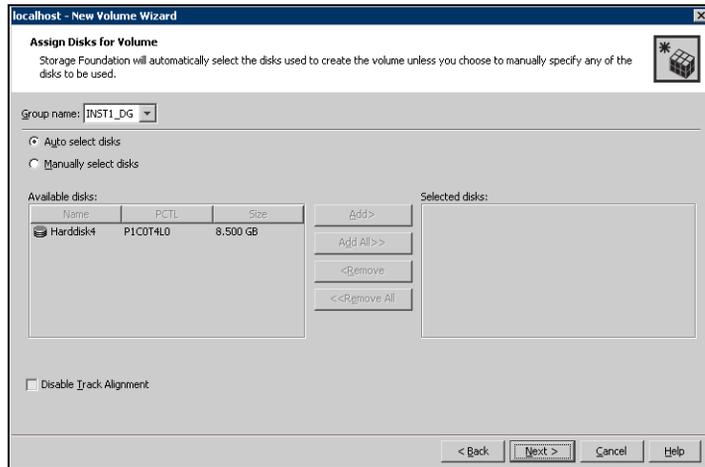
Warning: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

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

To create dynamic volumes

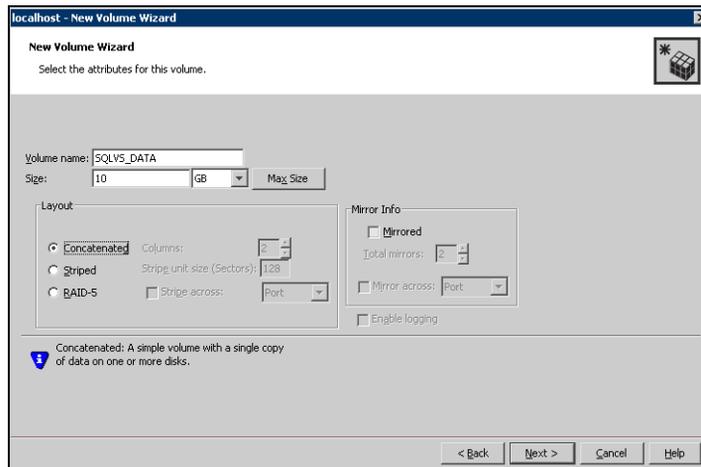
- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

6 Select the disks for the volume.



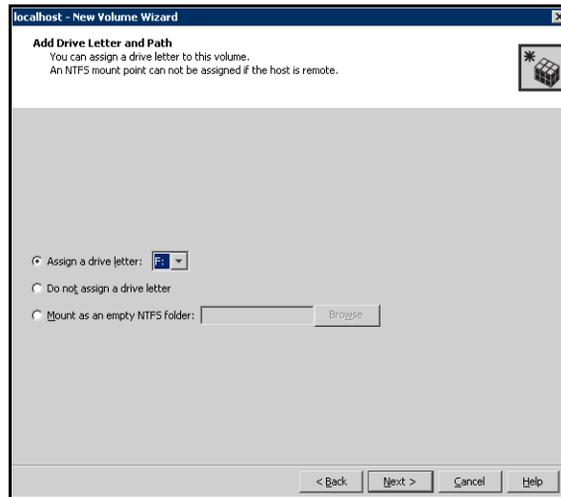
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.
- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.
 You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling **Track Alignment** means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the parameters of the volume.



- Enter a name for the volume. 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.
 - 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.
 - Verify that **Enable Logging** 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

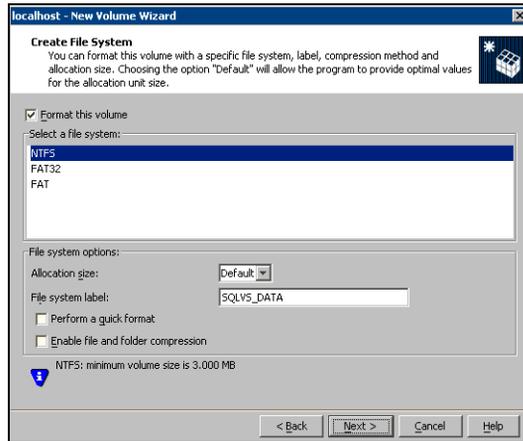
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- 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.
- For the Replicator Log volume only:
 Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - 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**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Installing and configuring SQL Server 2005 on the first node

Before installing Microsoft SQL Server 2005, ensure the cluster disk group is imported to the first node and the volumes are mounted.

Complete the following procedures to install and configure Microsoft SQL Server 2005:

- [Installing Microsoft SQL Server](#)
- [Setting the startup mode of the SQL Server 2005 services](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters). See “[Importing the cluster disk group](#)” on page 709 and “[Adding drive letters to mount the volumes](#)” on page 709.

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005

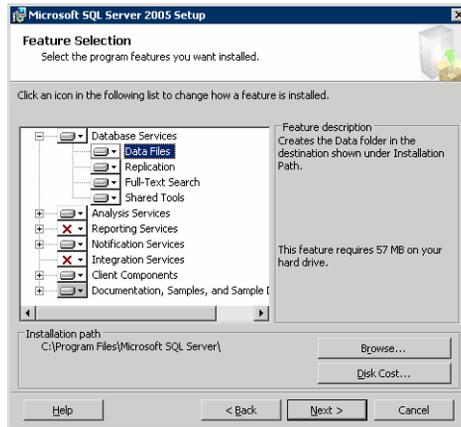
- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.

- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.

If you install optional components on one node, install the same components in the same order on other nodes.

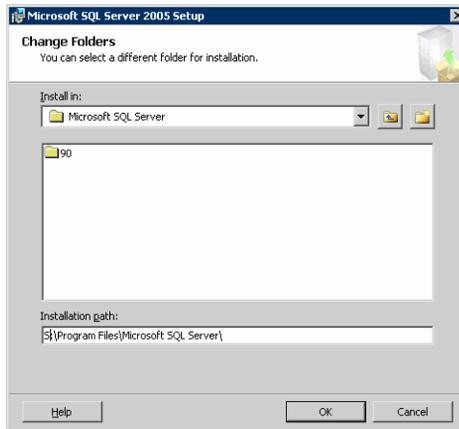
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.



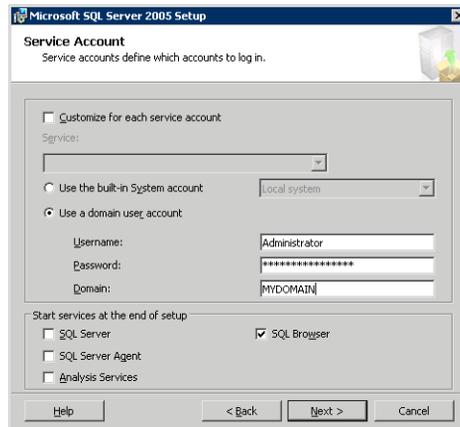
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 704, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

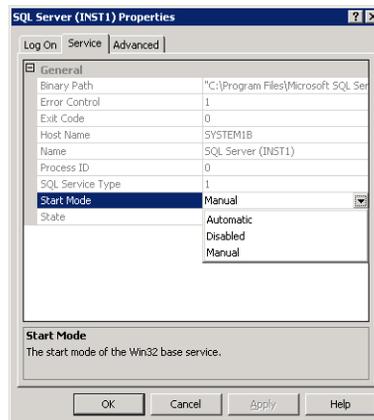
- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Setting the startup mode of the SQL Server 2005 services

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- “[Stopping the SQL Server 2005 Service](#)” on page 708
- “[Deporting the cluster disk group](#)” on page 708
- “[Importing the cluster disk group](#)” on page 709
- “[Adding drive letters to mount the volumes](#)” on page 709
- “[Renaming shared SQL Server 2005 files](#)” on page 710

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.

- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and if prompted, select a profile.
- 3 Click **Connect to a Host or Domain** and in the Connect dialog box, specify the host name and click **Connect**.
- 4 In the tree view, expand the system name (SYSTEM1), expand **Storage Agent**, and expand **Disk Groups**.
- 5 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 6 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) on the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the VEA, connect to the node where you want to import the cluster disk group.
- 2 In the tree view, expand the system name (SYSTEM2), right-click **Storage Agent**, and click **Rescan** to update the disk information on the node.
- 3 In the tree view, expand **Disk Groups**.
- 4 In the tree view, right-click the cluster disk group (INST1_DG) and select **Import Dynamic Disk Group**.
- 5 In the **Import Dynamic Disk Group** dialog box, click **OK**.

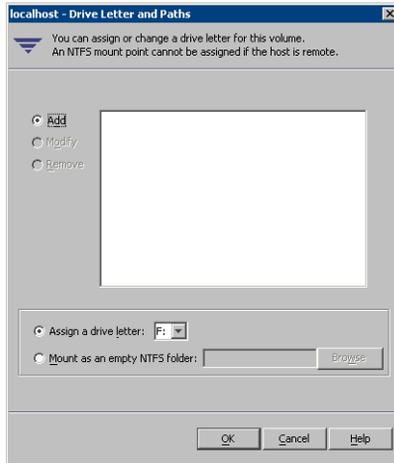
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL

Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 711
- [“Removing shared SQL Server files”](#) on page 715

Installing SQL Server

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters).

See [“Importing the cluster disk group”](#) on page 709 and [“Adding drive letters to mount the volumes”](#) on page 709.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

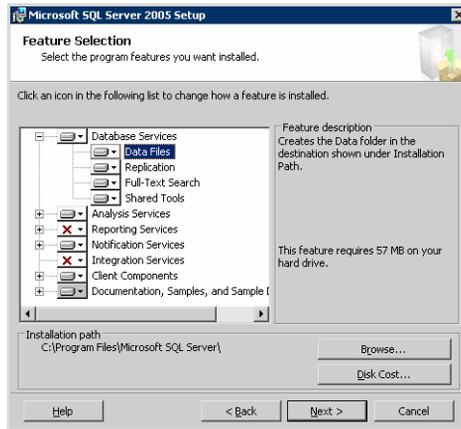
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005

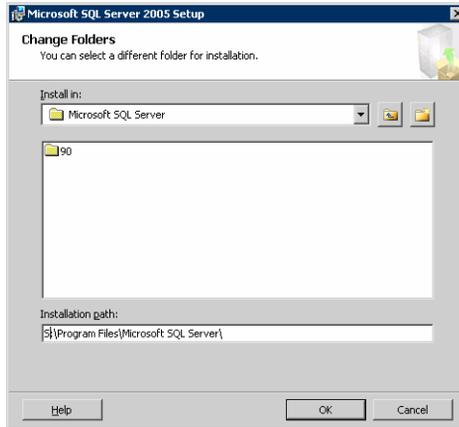
- 1 Navigate to the installation directory and launch **splash.hta**.

- 2 Under the Install section, click **Server components, tools, Books Online, and samples.**
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.



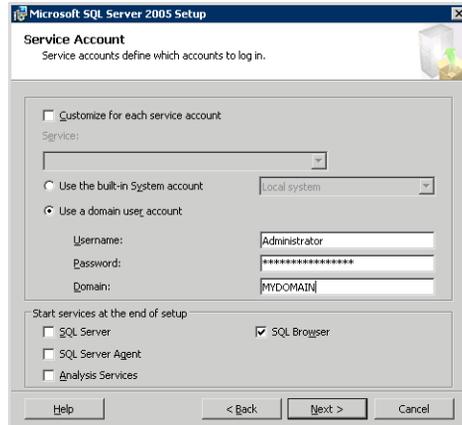
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 712, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
 Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Repeat the procedures described in “[Preparing to install SQL Server 2005 on the second node](#)” on page 708 and “[Installing SQL Server 2005 on the second node](#)” on page 711 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

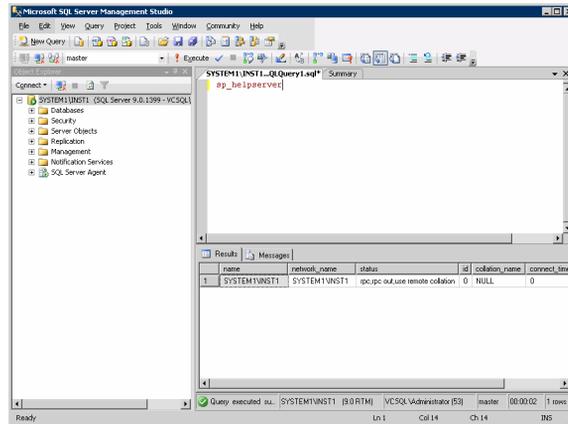
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name*\Instance_Name.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1\INST1. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
`sp_dropserver "System_Name\Instance_Name"`
 where *System_Name\Instance_Name* is the name noted in [step 3](#) on page 716.
 For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
 For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press **F5**.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.

Configuring the VCS SQL Server service group

The Enterprise Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Complete the following tasks before configuring the service group.

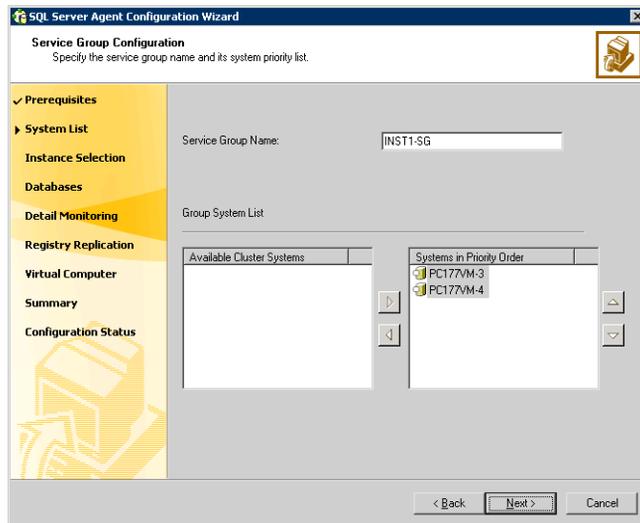
- Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2005, is installed on all cluster nodes.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See [“Configuring the cluster”](#) on page 678.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- Ensure that you are a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify that the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify that the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- Verify the virtual server name that was specified when setting the internal name of the clustered SQL Server instance. You specify this name when configuring the service group.
- Assign a unique virtual IP address to the SQL Server 2005 instance. You specify this IP address when configuring the service group.

- Optionally, to use a monitor script, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes.

A sample script is supplied in `C:\Program Files\Veritas\cluster server\bin\SQLServer2005\sample_script.sql`. Detailed monitoring is often not necessary.

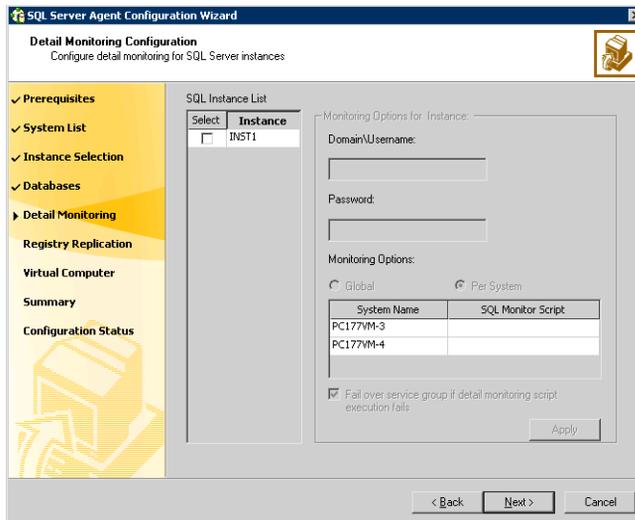
To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



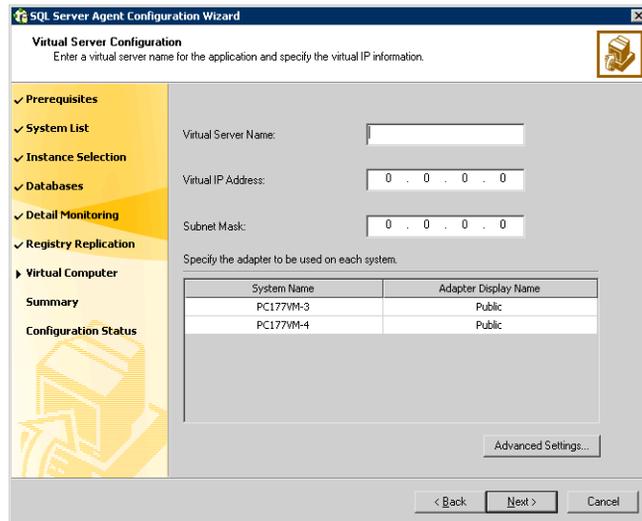
- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.

- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
Click **Next**.
 - 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
 - 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 720 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.

- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 720 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the

same as the virtual server name specified when setting the internal name of the clustered instance.

- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

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

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.

14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.

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

15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.

16 Complete the SQL Server service group configuration:

- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
- Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

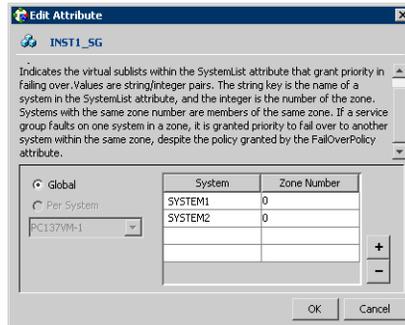
The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

Creating the primary system zone

In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the SQL Server service group (INST1_SG) in the left pane and the Properties tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone. Make sure you specify the systems in uppercase.



- 7 Click **OK**.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 697 for information on how to use the VEA console to create a volume.

Note: Best practice is to create a separate disk group with separate volumes for the user-defined database files.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.

- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as **CRITICAL**.

- Click **Next** to configure another SQL service group or an MSDTC service group.

Verifying the installation in the primary zone

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

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

To shut down an active cluster node

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

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

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, use the guidelines in the following list to complete the same tasks in the secondary zone.

Before you begin to configure the secondary zone, offline the following resources in the SQL service group in the primary zone:

- SQL Server resource (<sqlservicegroupname> - SQLServer2005)
- SQL Virtual Server name resource (<sqlservicegroupname> - Lanman)
- SQL Virtual IP resource (<sqlservicegroupname> - IP)

The remaining resources should be online, including the VMDg resources and the MountV resources.

In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

- [“Reviewing the prerequisites”](#) on page 659
- [“Reviewing the configuration”](#) on page 664
- [“Configuring the storage hardware and network”](#) on page 666
- [“Installing Veritas Storage Foundation HA for Windows”](#) on page 668
- [“Configuring VxSAS”](#) on page 675
- [“Adding the systems in the secondary zone to the cluster”](#) on page 728
- [“Configuring cluster disk groups and volumes”](#) on page 694

During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:

- Cluster disk group name
- Volume sizes
- Volume names
- Drive letters
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 703
Select the same options in the secondary zone as you did for the primary zone.
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 708

The instance name must be the same in the primary zone and secondary zone.

- [“Installing SQL Server 2005 on the second node”](#) on page 711

Note: After you install SQL Server 2005 on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes, to avoid conflicts during the configuration of the zones.

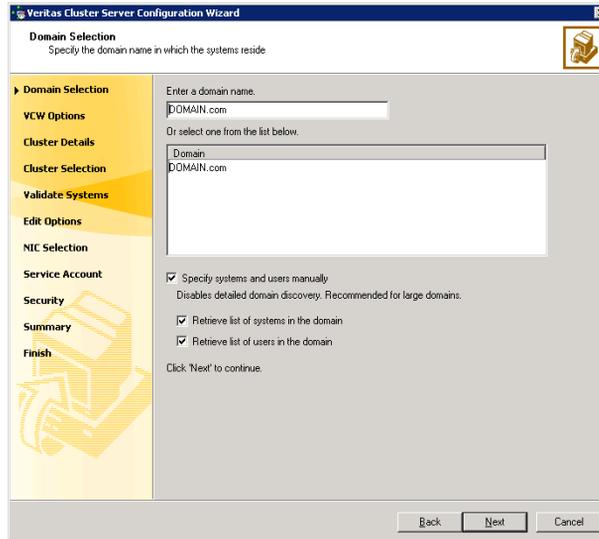
Adding the systems in the secondary zone to the cluster

Add the nodes in the secondary zone to the existing cluster with the following procedure.

To add a node to a VCS cluster

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

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



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

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

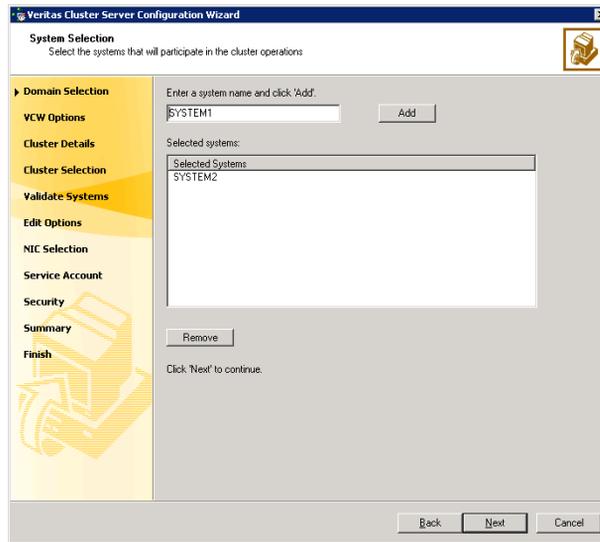
Proceed to [step 8](#) on page 732.

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

- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

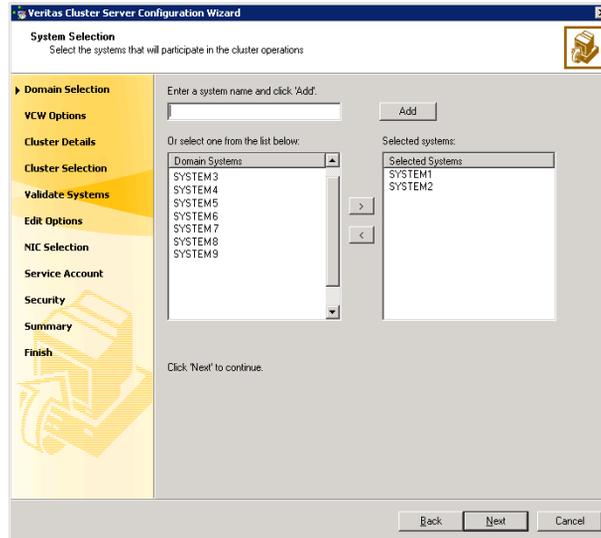
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 731. Otherwise proceed to the next step.

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



- Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
- If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 8](#) on page 732.

- 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.

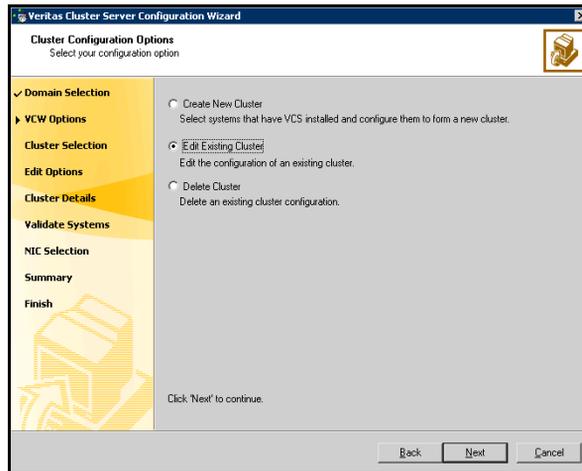
A system can be rejected for any of the following reasons:

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

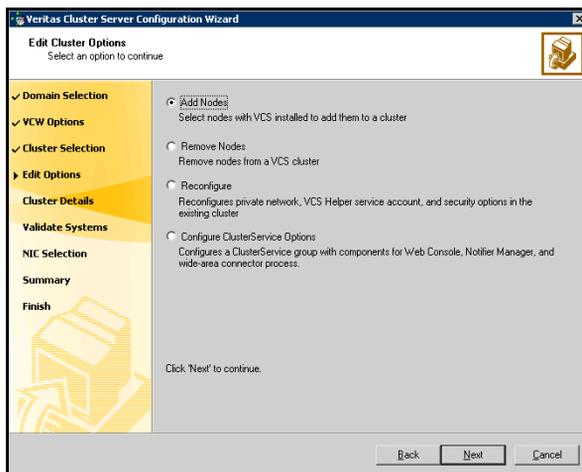
Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

Click **Next** to proceed.

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



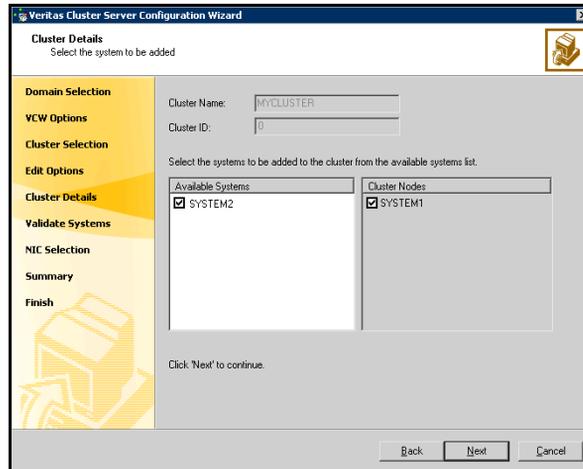
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**. If you chose to specify the systems manually in [step 4](#) on page 729, only the clusters configured with the specified systems are displayed.
- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **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.

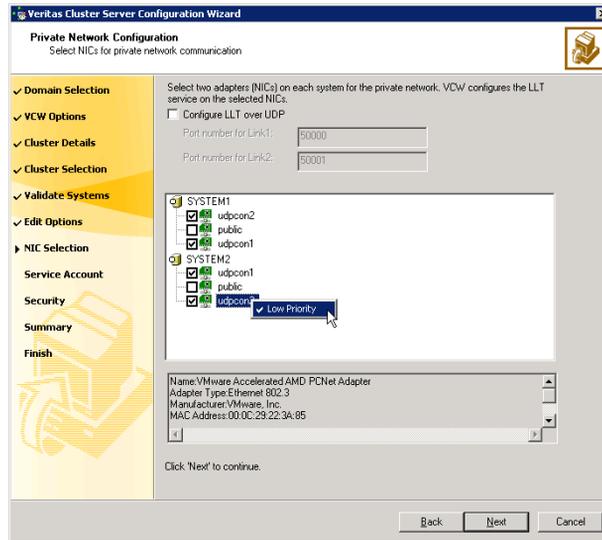
- 12 The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

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:

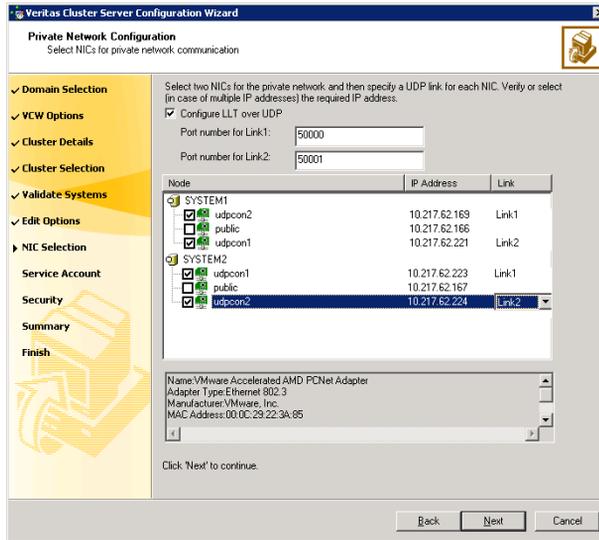
- To configure the VCS private network over Ethernet



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

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

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

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

- 15 Specify the 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**.

Return to the task list “[Creating a parallel environment in the secondary zone](#)” on page 727.

Setting up the Replicated Data Sets (RDS)

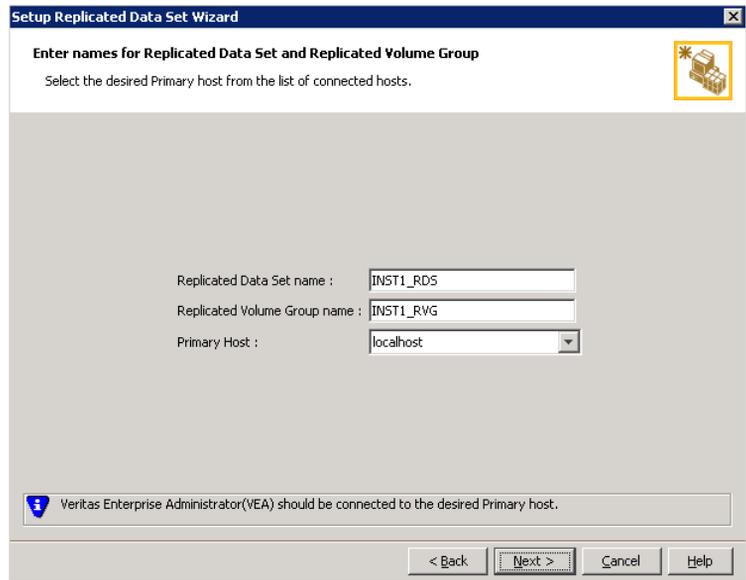
Set up the Replicated Data Sets (RDS) in the primary zone and secondary zone. You can configure an RDS using the Create RDS wizard for both zones.

- Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- Verify that the cluster disk group is imported and the volumes are mounted in the primary and secondary zone

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**. From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.

- 3 Read the Welcome page and click **Next**.
- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).

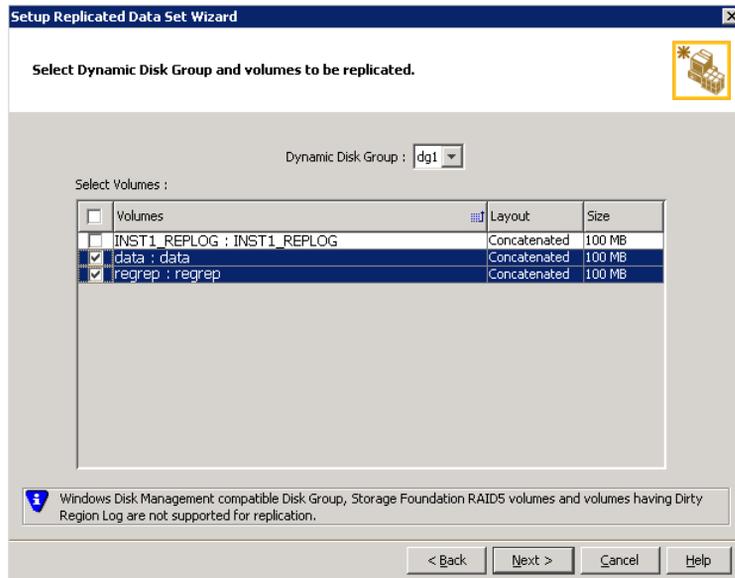


By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

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 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

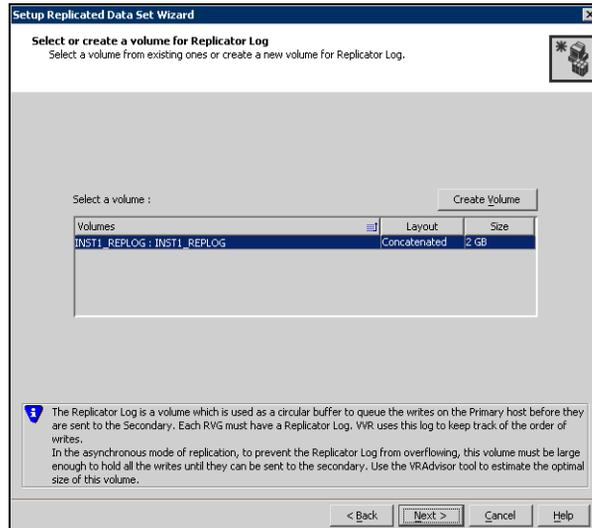


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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.

- 7 Click **Next**.

8 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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

Name Enter the name for the volume in the **Name** field.

Size Enter a size for the volume in the **Size** field.

Layout Select the desired volume layout.

- Disk Selection**
- Choose **Select disks automatically** if you want VVR to select the disks for the Replicator Log.
 - Choose **Select disks manually** to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select 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.
- 9 Review the information on the summary page and click **Create Primary RVG**.
 - 10 After the Primary RVG has been created successfully, VVR displays the following message:

RDS with Primary RVG has been created successfully. Do you want to add Secondary host to this RDS for replication now?
Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.
Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.
 - 11 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field. 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.
 - 12 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then VVR displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

 - the same or larger amount of space as that on the Primary
 - Enough space to create volumes with the same layout as on the Primary

Otherwise, the RDS setup wizard enables you to create the required volumes manually.

 - Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
 - Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.
 - 13 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.

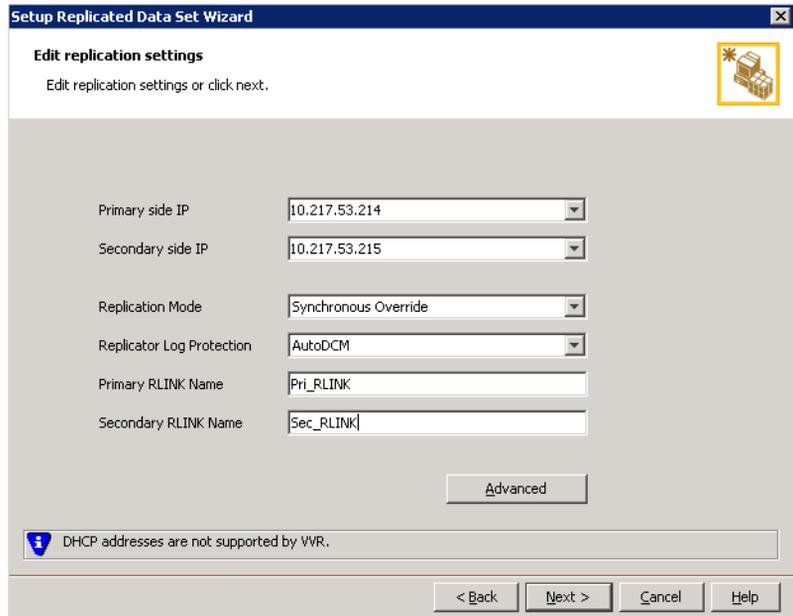
- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard.

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.

14 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, Synchronous, or Asynchronous . The default is synchronous override. Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous. Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates. Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously. If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.
Replicator Log Protection	The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection.

In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

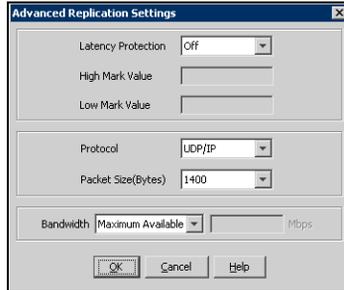
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- 15 Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

- **Off** is the default option and disables latency protection.
- **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
- **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize Automatically If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from Checkpoint If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 18 Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

If the SQL Server user-defined database and files are in a separate disk group from the SQL Server system files, repeat the procedure “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 736 for the disk group (INST1_DB1_DG) that contains the SQL Server user-defined database and files. Provide unique names for the Replicated Data Set name, and the Replicated Volume Group name. See “[Sample configuration](#)” on page 664 for a list of example names.

Configuring a hybrid RVG service group for replication

Create and configure a hybrid Replicated Volume Group (RVG) service group for replication.

The RVG service group is hybrid because it behaves as a failover service group within a zone and as a parallel service group between zones.

For additional information about service group types, see the *Veritas Cluster Server 5.1 Administrator's Guide*.

Configure the RVG service group's resources manually by copying and modifying components of the SQL Server service group. Then create new RVG resources and bring them online.

The RVG service group for RDC contains the following resources:

Table 14-5 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg for the first disk group	Volume Manager disk group with SQL system files
VvrRvg for the first disk group	Replicated volume group with SQL system files
VMDg for the second disk group	Volume Manager disk group with SQL user-defined files
VvrRvg for the second disk group	Replicated volume group with SQL user-defined files

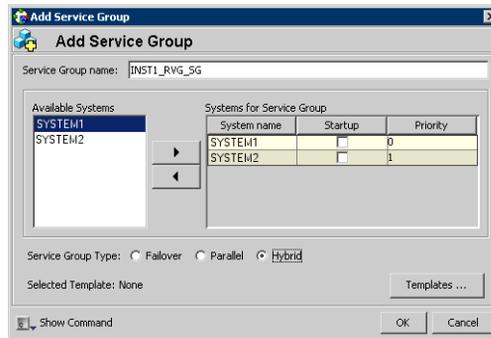
Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

To create a hybrid RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.

3 In the **Add Service Group** window:



- a Enter a name for the service group, for example INST1_RVG_SG. Make sure the service group name is in uppercase.
- b Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- c Select **Hybrid**.
- d Click **OK**.

Configuring the RVG service group for RDC replication

Configure the RVG service group's resources manually for RVG by completing the following tasks:

- [“Configuring the IP and NIC resources”](#)
Copy IP and NIC resources of the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Configuring the VMDg resources for the first disk group”](#) and [“Configuring the VMDg resources for the second disk group”](#)
Copy the VMDg resources for all disk groups in the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
- [“Adding the VVR RVG resources for the disk groups”](#)
Create the VVR RVG resources for all the disk groups and enter the attributes for each of the disk groups and the replication IP address.
- [“Linking the VVR RVG resources to establish dependencies”](#)
Link the VVR RVG resources to establish the dependencies between the VMDg resources, the IP resource for replication, and the VVR RVG resources for the disk groups. Configure the RVG service group's VMDg resources to point to the disk groups that contain the RVGs.

- “Deleting the VMDg resource from the SQL Server service group”
 Delete the VMDg resources from the SQL Server service group, because they depend on the replication and were configured in the RVG service group.

Configuring the IP and NIC resources

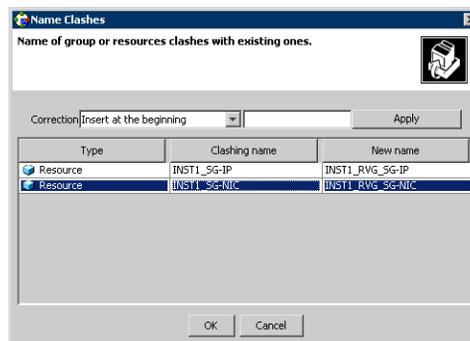
Configure the following resources and attributes for the IP and NIC:

Table 14-6 IP and NIC resources

Resource	Attributes to Modify
IP	Address
NIC	(none)

To create the IP resource and NIC resource

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the IP resource (INST1_SG-IP), and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 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.



- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the **Resources** tab display area, right-click the IP resource (INST1_RVG_SG-IP) 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 VVR IP address for the Primary Zone as the scalar value.
- 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 (INST1_RVG_SG-IP) and select **Enabled**.
- 2 In the **Resources** tab display area, right-click the NIC resource (INST1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg resources for the first disk group

Create the VMDg resource in the SQL Server service group, and clear the DGGuid attribute for the new VMDg.

Configure the following attributes in the SQL Server service group for the MountV resource:

Table 14-7 MountV resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
MountV (for the SQL Server system volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name
Resources for the disk group for the SQL user-defined database files:	
MountV (for the SQL Server user-defined database log)	VMDg Resource Name Volume Name
MountV (for the SQL Server user-defined database)	VMDg Resource Name Volume Name

To create the VMDg resource for the first disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the first disk group, with the SQL system files (INST1_SG-VMDg), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg.
- 6 Click **OK**.

To clear the DGGuid attribute for the new VMDg

- 1 In the **Resources** tab display area, right-click the new VMDg resource.
- 2 In the same **Properties View** window, for the **DGGuid** attribute, click **Edit**.
- 3 In the **Edit Attribute** window, clear the scalar value for the **DGGuid** attribute.
- 4 Close the **Properties View** window.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server system data files (INST1_SG-MountV) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server system data files (INST1_DATA_FILES).
- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the registry volume (INST1_SG-MountV-1) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the registry volume (INST1_REGREP_VOL).

- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg) and select **Enabled**.

Configuring the VMDg resources for the second disk group

Repeat the VMDg and MountV configuration for the second disk group.

To create the VMDg resource for the second disk group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the **Resources** tab, right-click the VMDg resource for the second disk group, with SQL user-defined files (INST1_SG-VMDg-1), and click **Copy > Self**.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.
- 5 In the **Name Clashes** window, change the name of the VMDg resource for the RVG service group, for example to INST1_RVG_SG-VMDg-1.
- 6 Click **OK**.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined log (INST1_SG-MountV-2) and select **View > Properties View**.
- 3 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined log (INST1_DB1_LOG).

- 4 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg-1.
- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined database (INST1_SG-MountV-3) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the **Volume Name** attribute is the SQL Server user-defined database (INST1_DB1_VOL).
- 9 In the same **Properties View** window, for the **VMDg Resource Name** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg resource that was just created, for example INST1_RVG_SG-VMDg.
- 11 Close the **Properties View** window.

To enable the VMDg resource

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the **Resources** tab display area, right-click the VMDg resource (INST1_RVG_SG-VMDg-1) and select **Enabled**.

Adding the VVR RVG resources for the disk groups

Add VVR RVG resources for replication of the disk groups.

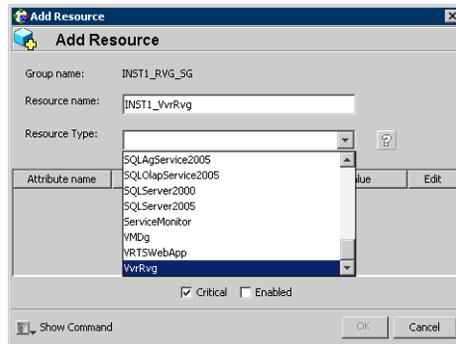
Configure the following attributes in the RVG service group for the VvrRvg resource:

Table 14-8 VvrRvg resources

Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
VvrRvg	VMDgResName IPResName
Resources for the disk group for the SQL user-defined database files:	
VvrRvg	VMDgResName IPResName

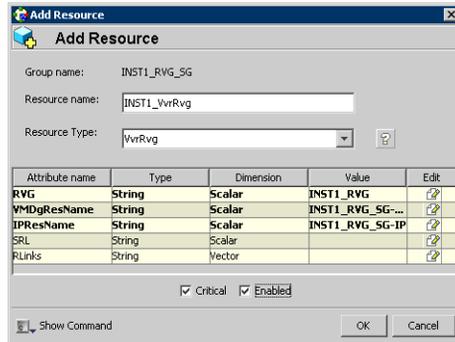
To create the VVR RVG resource for the first disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- 3 Enter the **Resource Name** for the VVR RVG resource.
- 4 Select the **Resource Type** of VvrRvg.
- 5 In the **Add Resource** window the attributes appear. For the **RVG** attribute, click **Edit**.
- 6 In the **Edit Attribute** window, enter the name of the RVG group that is being managed, for example INST1_RVG.
- 7 Click **OK**.
- 8 In the **Add Resource** window, for the **VMDGResName** attribute, click **Edit**.
- 9 In the **Edit Attribute** window, enter the name of the disk group containing the RVG, for example INST1_RVG_SG-VMDg.
- 10 Click **OK**.
- 11 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 12 In the **Edit Attribute** window, enter the name of the IP resource managing the IP address for replication, for example INST1_RVG_SG-IP.
- 13 Click **OK**.

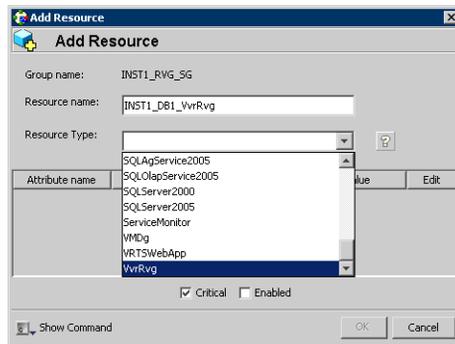
12 In the **Add Resource** window, verify that the attributes have been modified:



13 Click **OK**.

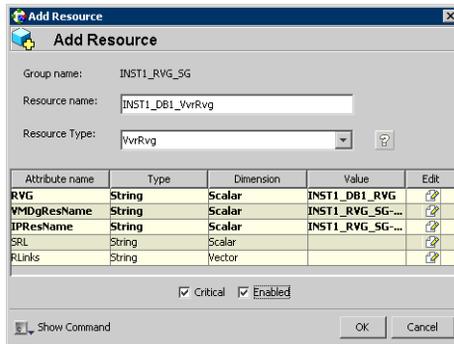
To create the VVR RVG resource for the second disk group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the **Add Resource** window:



- Enter the **Resource Name** for the VVR RVG resource for the second disk group.
 - 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 group that is being managed, for example INST1_DB1_RVG.
 - 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 disk group containing the RVG, for example INST1_RVG_SG-VMDg-1.
- 8 Click **OK**.
- 9 In the **Add Resource** window, for the **IPResName** attribute, click **Edit**.
- 10 In the **Edit Attribute** window, enter the name IP resource managing the IP address for replication, for example INST1_RVG_SG-IP. In this example both disk groups are using the same IP resource for replication.
- 11 Click **OK**.
- 12 In the **Add Resource** window, verify that the attributes have been modified:



- 13 Click **OK**.

Linking the VVR RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the VVR RVG service group to establish the dependencies between the resources. Start from the top parent and link the following resources:

Table 14-9 Dependencies for VVR RVG resources for RDC

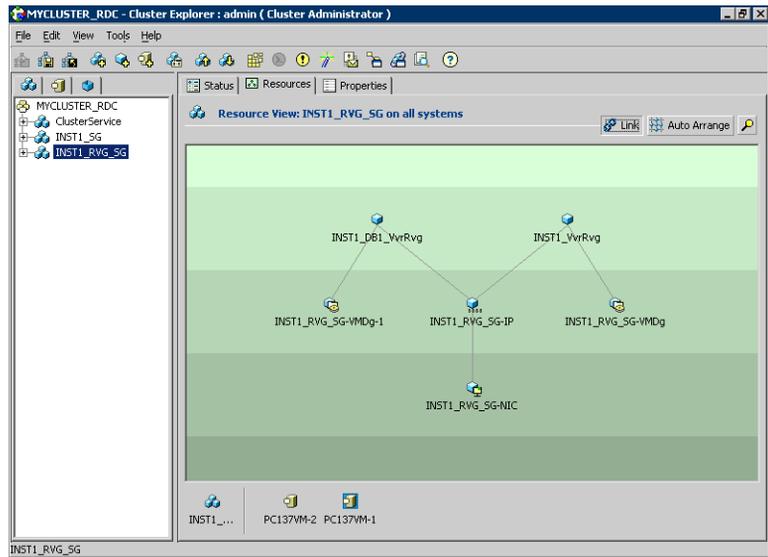
Parent	Child
Resources for the disk group for the SQL system files:	
INST1_VvrRvg	The IP for replication, for example INST1_RVG_SG-IP
INST1_VvrRvg	The VMDg for the SQL system files, for example INST1_RVG_SG-VMDg
Resources for the disk group for the SQL user-defined database files:	
INST1_DB1_VvrRvg	The IP for replication, for example INST1_RVG_SG-IP

Table 14-9 Dependencies for VVR RVG resources for RDC

Parent	Child
INST1_DB1_VvrRvg	The VMDg for the SQL user-defined database files, for example INST1_RVG_SG-VMDg-1

To link the VVR RVG resources

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource, for example INST1_DB1_VvrRvg.
- 4 Click the child resource, for example INST1_RVG_SG-IP.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG resources:



Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg resource from the SQL Server service group

The VMDg resources must now be manually deleted from the SQL Server service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg Resources from the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) from the left pane.
- 2 In the **Resources** tab display area, right-click the VMDg resource for the first disk group (INST1_SG-VMDg) and select **Delete**.
- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the **Resources** tab display area, right-click the VMDg resource for the second disk group (INST1_SG-VMDg-1) and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVGPrimary to the SQL Server service group for each of the SQL Server disk groups (system and user-defined) and configure the attributes.

Set the value of the **RvgResourceName** attribute to the name of the RVG resource for the RVGPrimary agent.

Configure the following attributes in the SQL Server service group for the RVG Primary resources:

Table 14-10 RVG Primary resources

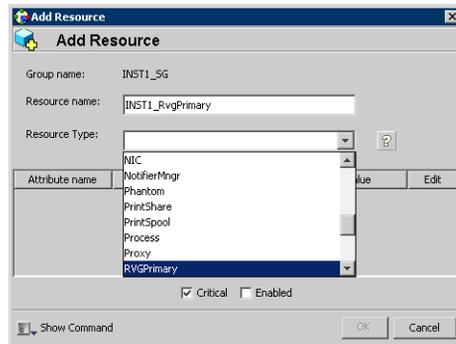
Resource	Attributes to Modify
Resources for the disk group for the SQL system files:	
RVGPrimary	RvgResourceName
Resources for the disk group for the SQL user-defined database files:	
RVGPrimary	RvgResourceName

Creating the RVG Primary resources

For all disk groups, create an RVG Primary Resource for replication.

To create the RVG Primary resource for the SQL Server system disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:



- Enter the **Resource Name** for the RVG Primary resource for the SQL Server system files disk group, for example INST1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
 - 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_VvrRvg and click **OK**.
 - 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Veritas Cluster Server 5.1 Administrator's Guide* for more information about the RVG Primary agent.
 - 6 Verify that **Critical** and **Enabled** are both checked.
 - 7 Click **OK**.

To create the RVG Primary resource for the SQL Server user-defined database disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window:
 - Enter the **Resource Name** for the RVG Primary resource for the SQL Server user-defined database disk group, for example INST1_DB1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the **Add Resource** window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
- 4 In the **Edit Attribute** window, enter the name of the VVR RVG resource, for example INST1_DB1_VvrRvg and click **OK**.
- 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.
- 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 SQL Server service group (INST1_SG) to establish the dependencies between the resources for replication.

Start from the top parent and link the following resources:

Table 14-11 Dependencies for the RVG Primary resources for RDC

Parent	Child
INST1_SG-MountV	INST1_RvgPrimary
INST1_SG-MountV-1	INST1_RvgPrimary
INST1_SG-MountV-2	INST1_DB1_RvgPrimary
INST1_SG-MountV-3	INST1_DB1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 Click the **Link** button in the right pane.

- 3 Click the parent resource, for example INST1_SG-MountV.
- 4 Click the child resource, for example INST1_RvgPrimary.
- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the SQL Server service group (INST1_SG) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane on the **Resources** tab, right-click the first RVG Primary resource (INST1_RvgPrimary) and select **Online > SYSTEM1**.
- 3 In the right pane on the **Resources** tab, right click the second RVG Primary resource (INST1_DB1_RvgPrimary) and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG

In the RVG service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG Service Group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone.
- 7 Click **OK**.

Setting a dependency between the service groups

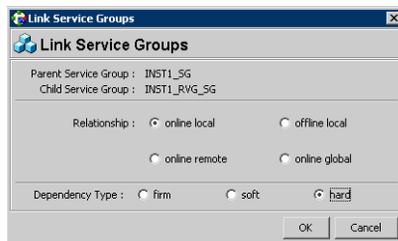
The RVG service group must be online on both the primary and secondary zones. However, if a failover occurs from one node to another within the same zone, the RVG service group must fail over along with the application service group.

To ensure that the SQL Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the SQL Server service group.

The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

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 SQL Server service group (the parent service group), for example INST1_SG.
- 5 Click the RVG service group (the child resource), for example INST1_RVG_SG.
- 6 In the **Link Service Groups** window:



- Select the **Relationship** of **online local**.
- Select the **Dependency Type** of **hard**.
- Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

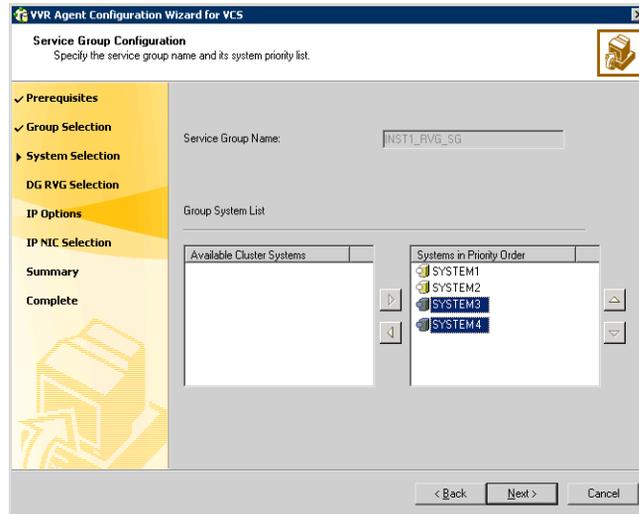
Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

To add the nodes from the secondary zone to the RVG

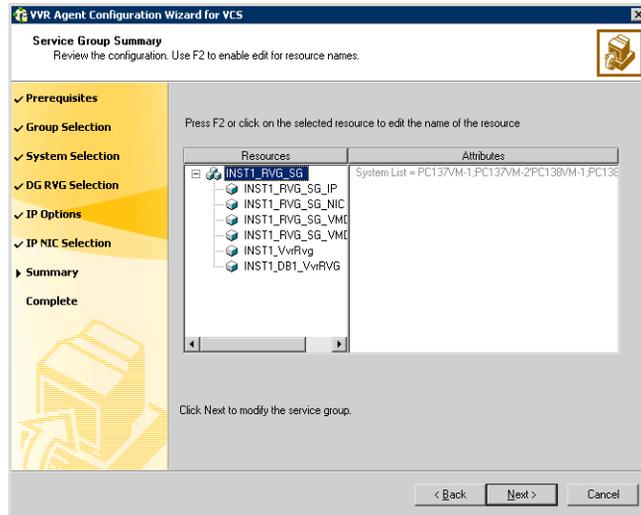
- 1 From the active node of the cluster in the primary zone, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Read and verify the requirements on the **Welcome page**, and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - 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:



- In the **Available Cluster Systems** box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 - To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
 - Click **Next**.
- 6 If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
 - 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
 - 8 In the IP Resource Options dialog box, select **Modify IP resource** and click **Next**.
 - 9 If a VCS error appears, click **OK**.
 - 10 In the Network Configuration dialog box, verify that the selected adapters are correct and click **Next**.

11 Review the summary of the service group configuration:



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

- Click **Next** to modify the replication service group.

12 When prompted, click **Yes** to modify the service group.

13 Click **Finish**.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the IP resources for failover

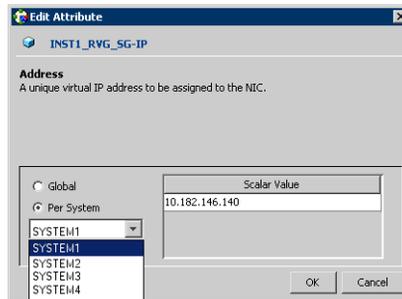
Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.

In the event of a system or SQL 2005 failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

To modify the IP resources in the RVG service group

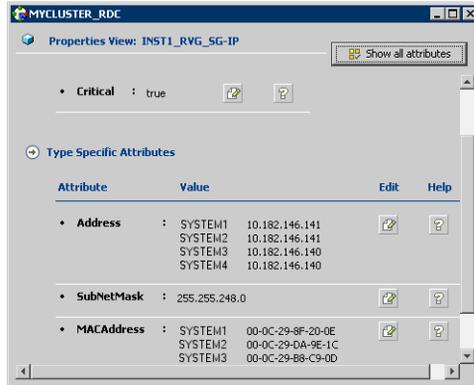
- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.

- 4 In the Edit Attributes window, edit the Address attribute.



- 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 (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5 In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone

have the same IP address, that is different from the IP address for the primary zone.



- 6 Close the Properties View window.
- 7 Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding the nodes from the secondary zone to the SQL Server service group

Use the SQL Server Agent Configuration Wizard to add the nodes from the secondary zone to the SQL Server service group.

To add the nodes from the secondary zone to the SQL Server service group

- 1 Select **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Configuration Wizard**.
- 2 In the Select Configuration Option dialog box, select **MS SQL Server - Service Group Configuration**. Select **Edit**, and click **Next**.
- 3 Verify that you have met the prerequisites listed and click **Next**.
- 4 Select the service group to be modified (INST1_SG) and click **Next**.
- 5 If a VCS notice message appears indicating that resources are online, click **Yes** to continue.
- 6 On the Service Group Configuration page, select the nodes in the secondary zone, use the arrow button to move them from **Available Cluster Systems** to **System in Priority Order**.

To change the priority of a system in the **Systems in Priority Order** list, select the system and click the up and down arrow icons. Arrange the

systems in priority order in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online.

This set of nodes selected for the SQL Server 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.

- 7 Click **Next**.
- 8 On the SQL Server Instance Selection page, click **Next**.
- 9 The User Databases List page summarizes the databases for this instance of SQL. Click **Next**.
- 10 On the Detail Monitoring Configuration page, clear the box in the SQL Instance List to disable monitoring, as required. Detailed monitoring is not necessary. Click **Next**.
- 11 On the Registry Replication Path page, click **Next**.
- 12 On the Virtual Server Configuration page, verify that the public adapter is used on each system and click **Next**.
- 13 In the Service Group Summary, review the service group configuration and click **Next**.
- 14 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.
- 15 Click **Finish**.

Configuring the zones in the SQL Server service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the SQL Server service group

- 1 From VCS Cluster Explorer, in the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.

- 7 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- [Bringing the service group online](#)
- [Switching online nodes](#)

Bringing the service group online

After completing all configuration, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the Exchange Server service group online in the primary zone.

To bring the Exchange service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the Exchange Server service group (EVS1_SG1).
- 2 Click **Online**.

Switching online nodes

Failover simulation is an important part of configuration testing. Test the failover by switching 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) (**Start > All Programs > Veritas > Veritas Cluster Manager (Java Console)**).
- 2 Click **Click here to log in** for the appropriate cluster. If this is your first use of the Veritas Cluster Manager, in the File menu, click **New Cluster**. In the **New Cluster - Connectivity Configuration** window, enter the computer name in the **Host name** field and click **OK**.
- 3 In the **Machinename - Login window**, enter your user name and password in the respective fields and click **OK**.
- 4 Right-click the service group in the left pane, and select an alternate system name from the **Switch To** entry.
- 5 In the **Question** dialog box, click **Yes** to confirm you do want to switch the service group to the other node.

Additional instructions for GCO disaster recovery

After completing the tasks for setting up a replicated data cluster for SQL Server 2005, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, if a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment. The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up VVR replication for the secondary site
- Configures the primary and secondary site clusters as global clusters

See [Chapter 17, “Deploying disaster recovery: New SQL Server 2005 installation”](#) on page 881.

Disaster Recovery

This section includes the following chapters:

- [Disaster recovery for SQL: Overview](#)
- [Deploying disaster recovery: New SQL Server 2000 installation](#)
- [Deploying disaster recovery: New SQL Server 2005 installation](#)
- [Configuring an MSDTC service group for disaster recovery](#)
- [Testing fault readiness by running a fire drill](#)

Disaster recovery for SQL: Overview

This chapter includes the following topics:

- [“What is a disaster recovery solution?”](#) on page 775
- [“What needs to be protected in a SQL Server environment?”](#) on page 777
- [“Typical disaster recovery configuration”](#) on page 777

What is a disaster recovery solution?

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.

A disaster recovery (DR) solution is a series of procedures which you can use to safely and efficiently restore application user data and services in the event of a catastrophic failure. A typical DR solution requires that you have a source or *primary site* and a destination or *secondary site*. The user application data on the primary site is replicated to the secondary site. The cluster on the primary site provides data and services during normal operations; in the event of a disaster at the primary site and failure of the cluster, the secondary site provides the data and services.

Why implement a disaster recovery solution?

A DR solution is vital for businesses that rely on the availability of data. A well-designed DR solution prepares a business for unexpected disasters and provides the following benefits in a DR situation:

- Minimizes economic loss due to the unavailability or loss of data.
- Provides a plan for the safe and orderly recovery of data in the event of a disaster.
- Ensures safe and efficient recovery of data and services.
- Minimizes any decision making during DR.
- Reduces the reliance on key individuals.

Strategically planning a DR solution provides businesses with affordable ways to meet their service level agreements, comply with government regulations, and minimize their business risks.

Note: A DR solution requires a well-defined backup strategy. Refer to Veritas NetBackup product documentation for information on configuring backup.

Understanding replication

The term replication refers to the use of a tool or service to automate the process of maintaining a consistent copy of data from a designated source (primary site) on one or more remote locations (secondary sites).

In the event that the primary site data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site.

SFW HA provides Veritas Volume Replicator for use in replication. The SFW HA disaster recovery solution also supports hardware replication.

For more information on VVR refer to the *Veritas Volume Replicator Administrator's Guide*.

What needs to be protected in a SQL Server environment?

The following components of a SQL Server environment must be protected in the event of a disaster:

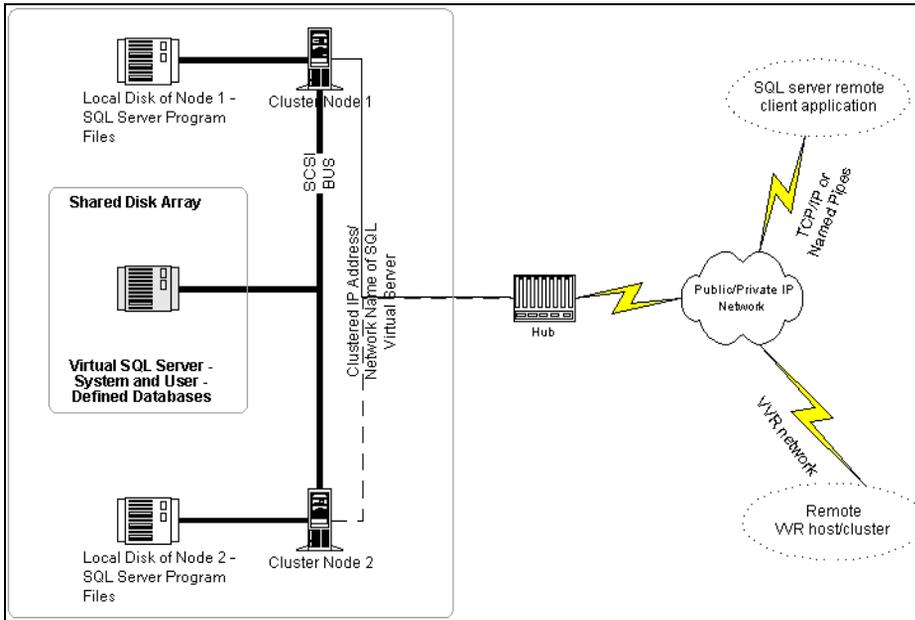
- **User Databases:** The most critical component in any SQL Server implementation is the user data that is stored in user-defined databases.
- **Logins:** Logins allow clients to connect to SQL Server and execute queries on user data. Logins are stored in the master database and each of the user-defined databases.
- **Jobs:** Jobs are a set of scheduled tasks that maintain SQL Server databases. The job configuration is stored in the msdb system database.
- **Alerts:** Alerts are actions that are taken when a specific event occurs. They are used to respond to and correct errors that occur in SQL Server. The alert configuration is stored in the msdb system database.
- **Operators:** Operators are contacts that address problems occurring in SQL Server. They are notified in the event of errors. The operator configuration is stored in the msdb system database.
- **Extended Stored Procedures:** Extended stored procedures are external routines that are called from within SQL Server. They are typically stored in DLL files on the file system.
- **Other Server Extensions:** SQL Server is a very flexible database engine and it is possible to extend its functionality in several ways. These extensions are also important to the operation of the SQL Server.

Typical disaster recovery configuration

The illustration below shows a typical configuration used to protect SQL Server in a clustered disaster recovery environment.

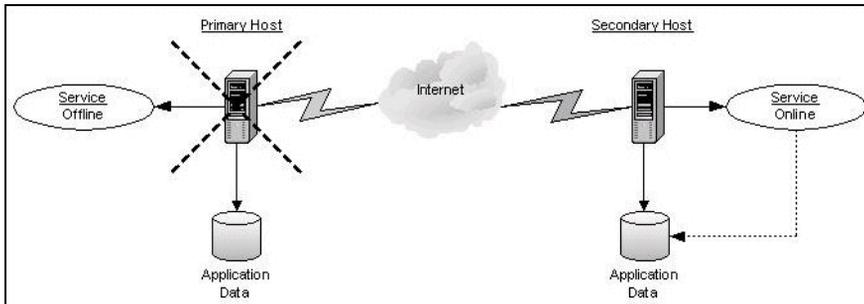
The VVR primary site is replicating the SQL Server application data to the secondary site. The setup contains clusters at the primary and secondary sites to make both SQL Server and VVR highly available.

Figure 15-1 Disaster recovery environment



When a disaster such as an earthquake causes a failure at the primary site, the secondary site takes over the primary role and restores the application services and data to users.

Figure 15-2 Secondary site online after a disaster



Deploying disaster recovery: New SQL Server 2000 installation

This chapter covers the following topics:

- [Tasks for a new disaster recovery installation of Microsoft SQL Server 2000](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Setting up the secondary site: Configuring SFW HA and setting up a cluster](#)
- [Verifying your primary site configuration](#)
- [Setting up your replication environment](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)](#)
- [Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)](#)
- [Installing and configuring SQL Server 2000 on the first node \(secondary site\)](#)
- [Preparing to install SQL Server on the second node \(secondary site\)](#)
- [Installing SQL Server 2000 on the second node \(secondary site\)](#)

- [Cloning the service group configuration from the primary to the secondary site](#)
- [Configuring replication and global clustering](#)
- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Adding multiple DR sites \(optional\)](#)
- [Recovery procedures for service group dependencies](#)

Tasks for a new disaster recovery installation of Microsoft SQL Server 2000

Before setting up SFW HA disaster recovery at a secondary site, you must complete the high availability configuration on the primary site.

See [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#) on page 49.

You can also use configure disaster recovery for a primary site that is configured as a replicated data cluster.

See [Chapter 13, “Configuring Replicated Data Clusters for SQL 2000”](#) on page 539.

After setting up an SFW HA environment for SQL on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Note: If you want to create the DR site configuration manually, without using the DR wizard, see [Appendix A, “Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation”](#).

To configure MSDTC service groups, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 1171.

[Table 16-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 16-1 Tasks for deploying disaster recovery for SQL Server 2000

Objective	Tasks
“Reviewing the requirements” on page 783	Verifying hardware and software prerequisites
Reviewing the configuration	<ul style="list-style-type: none"> ■ Understanding site failover in a DR environment ■ Reviewing the sample configuration ■ Understanding supported disaster recovery configurations for service group dependencies
Configuring the storage hardware and network	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Setting up the secondary site: Configuring SFW HA and setting up a cluster” on page 794	<ul style="list-style-type: none"> ■ Installing SFW HA ■ Configuring the cluster using the VCS Cluster Configuration Wizard (VCW)
“Verifying your primary site configuration” on page 817	Verifying that SQL has been configured for high availability at the primary site and that the service groups are online
“Setting up your replication environment” on page 817	Ensuring that replication prerequisites for your selected method of replication are met before running the DR wizard
“Assigning user privileges (secure clusters only)” on page 825	Assigning user privileges for a secure cluster

Table 16-1 Tasks for deploying disaster recovery for SQL Server 2000

Objective	Tasks
“Configuring disaster recovery with the DR wizard” on page 826	<ul style="list-style-type: none"> ■ Reviewing the prerequisites for the DR wizard ■ Starting the DR wizard and making the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method
“Cloning the storage on the secondary site using the DR wizard (VVR replication option)” on page 831	(VVR replication option) Cloning the storage configuration on the secondary site using the DR wizard
“Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 835	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Using the DR wizard to create temporary storage for installation on the secondary site
“Installing and configuring SQL Server 2000 on the first node (secondary site)” on page 839	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Configuring SQL services
“Preparing to install SQL Server on the second node (secondary site)” on page 842	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters or paths as needed ■ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2000 on the second node (secondary site)” on page 846	Installing SQL Server 2000
“Setting the internal name of the clustered instance” on page 850	Setting the internal name of the clustered instance
“Cloning the service group configuration from the primary to the secondary site” on page 852	Cloning the service group configuration from the primary to the secondary site using the DR wizard

Table 16-1 Tasks for deploying disaster recovery for SQL Server 2000

Objective	Tasks
“Configuring replication and global clustering” on page 856	<ul style="list-style-type: none"> ■ (VVR replication) Using the wizard to configure replication and global clustering ■ (EMC SRDF replication) Setting up replication and then using the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Setting up replication and then using the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Using the wizard to configure global clustering, and then setting up replication
“Verifying the disaster recovery configuration” on page 872	Verifying that the secondary site has been fully configured for disaster recovery
“Establishing secure communication within the global cluster (optional)” on page 874	Adding secure communication between local clusters within the global cluster (optional task)
“Adding multiple DR sites (optional)” on page 876	Optionally, adding additional DR sites to a VVR environment
“Recovery procedures for service group dependencies” on page 877	Bringing the service groups online after failover to the secondary site

Reviewing the requirements

This DR solution requires a primary site and secondary site.

Review the following installation and configuration requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

- “Disk space requirements” on page 784
- “Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)” on page 784

Disk space requirements

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

[Table 16-2](#) on page 784 estimates disk space requirements for SFW HA.

Table 16-2 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at:

<http://www.symantec.com/business/support/index.jsp>

Supported Software

- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the Global Clustering Option and depending on your replication solution, Veritas Volume Replicator or a hardware replication agent

For a Microsoft SQL Server 2000 environment, any of the following SQL Servers and their operating systems:

- | | |
|---|---|
| Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required) | ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) |
| | ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) |
| Microsoft SQL Server 2000 (64-bit) Enterprise Edition | ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |

- Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions)
- Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions)

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 787.
- 1 GB of RAM per server for SFW HA.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL Virtual Server
 - One IP address for each physical node in the cluster
 - One static IP address per cluster used when configuring the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *VCS Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Reviewing the requirements”](#) on page 783.

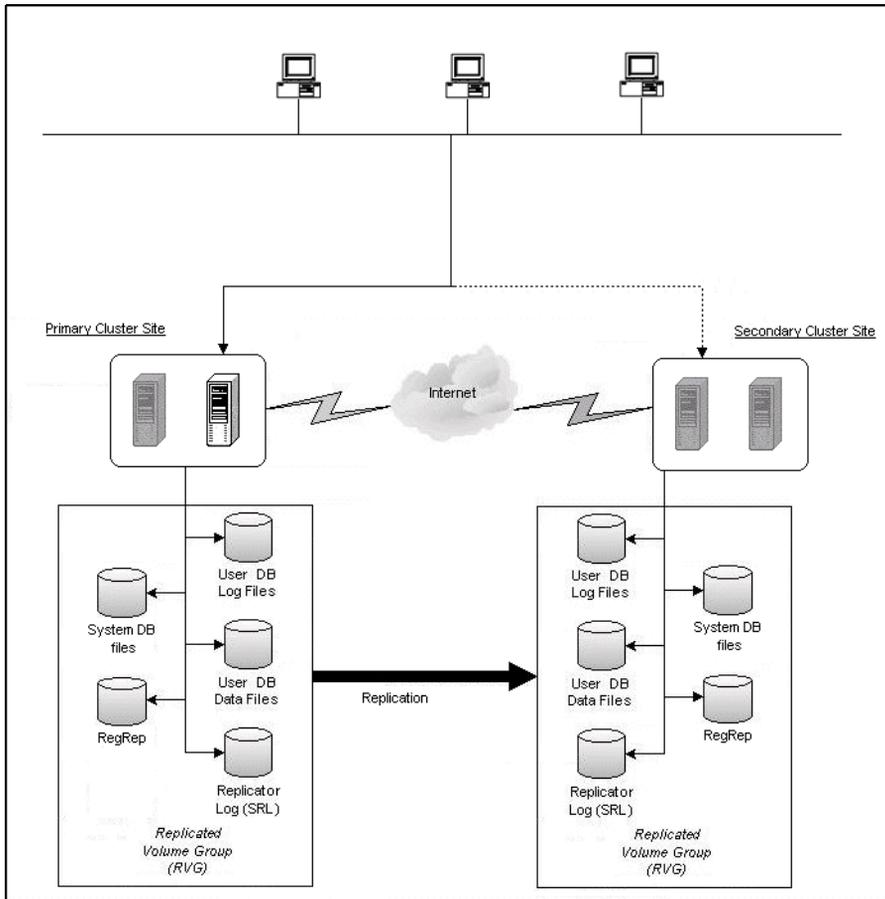
You should have these IP addresses available before you start deploying your environment.

The following figure illustrates a typical clustered VVR configuration for an active/passive configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2000 application data is stored on the volumes that are under the control of the RVG.

You can also configure disaster recovery for an active/active configuration.

See “[Active-Active configuration](#)” on page 62 in [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#).

Figure 16-1 Typical VVR configuration



If the Microsoft SQL Server 2000 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

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:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	(VVR only) replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
----------------------	--

All the other parameters are the same as on the primary site.

DR Components (VVR only)

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Setting up the secondary site: Configuring SFW HA and setting up a cluster

Install SFW HA and configure the cluster at the secondary site.

Begin with verifying that the requirements are met on the secondary site:

- See “[Reviewing the requirements](#)” on page 783.

Then continue with the following topics:

- “[Installing SFW HA](#)” on page 794
- “[Configuring the cluster](#)” on page 801

Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 16-3](#) on page 794 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 16-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed

Table 16-3 Installation behavior with unsigned drivers (Continued)

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 795.

To change the driver signing options on each system

- 1 Log on locally to the system.
 - 2 Open the Control Panel and click **System**.
 - 3 Click the **Hardware** tab and click **Driver Signing**.
 - 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
 - 5 Click **OK**.
 - 6 Repeat for each computer.
- If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

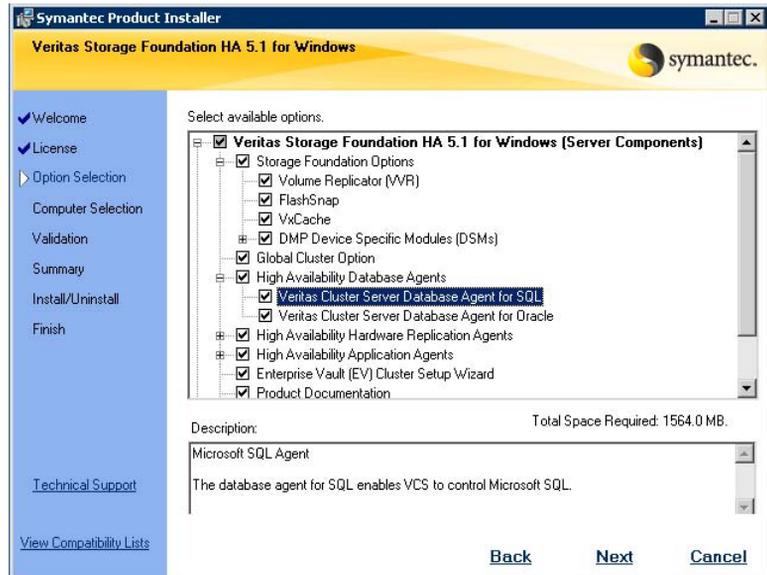
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

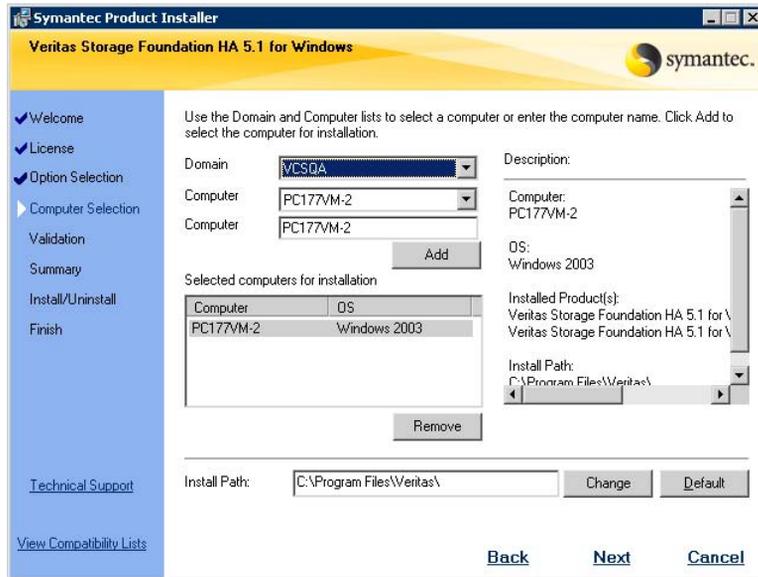
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



- | | |
|---|---|
| Veritas Cluster Server Database Agent for SQL | Required to configure high availability for SQL Server. |
| Client | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Veritas Volume Replicator | For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR. |
| High Availability Hardware Replication Agents | If you plan to use hardware replication, select the appropriate hardware replication agent. |

9 Select the domain and the computers for the installation and click **Next**.



- Domain** Select a domain from the list.
Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.
- Computer** To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.
To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.
Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
 If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
 If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

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

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

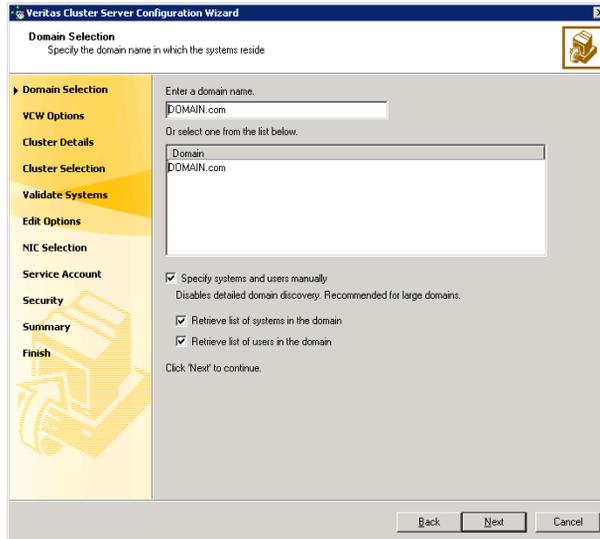
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



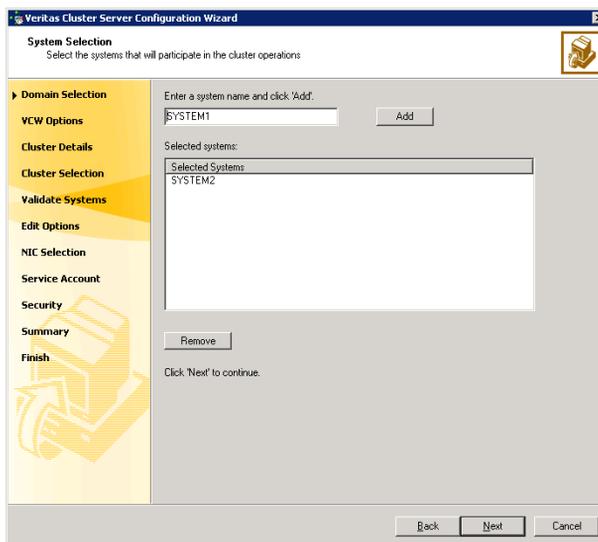
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
Proceed to [step 8](#) on page 804.

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

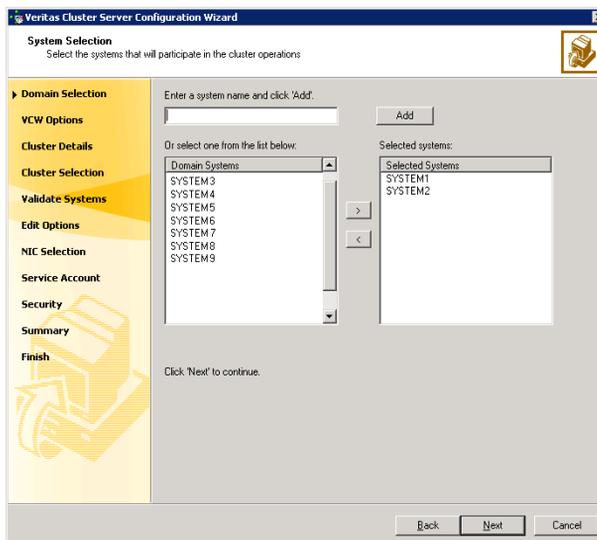
- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 803. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 804.

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



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

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

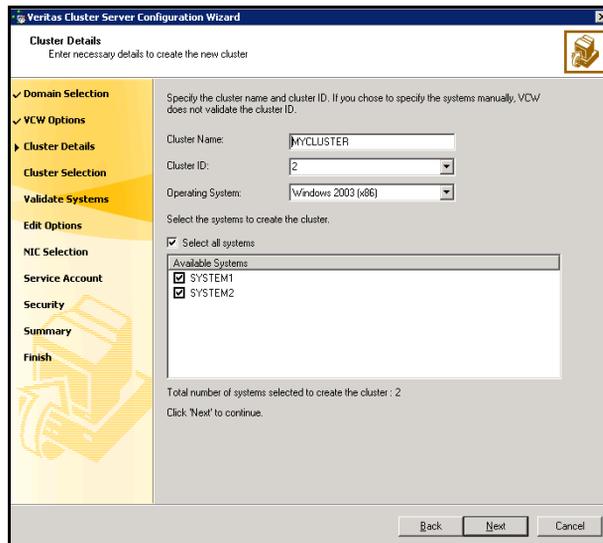
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 802 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

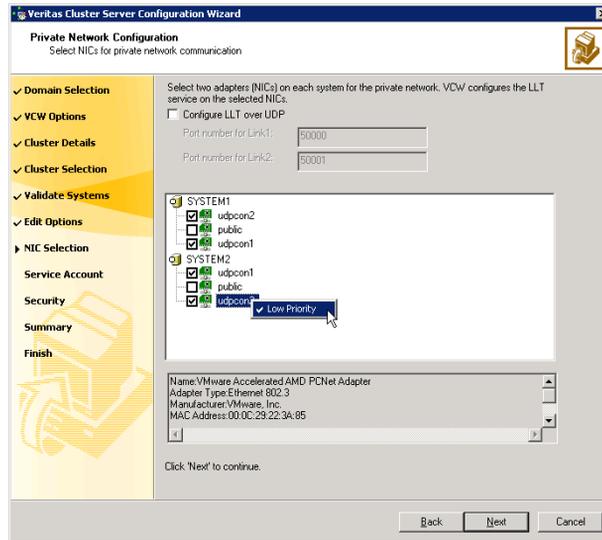
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 804, proceed to the next step. Otherwise, proceed to [step 12](#) on page 808.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

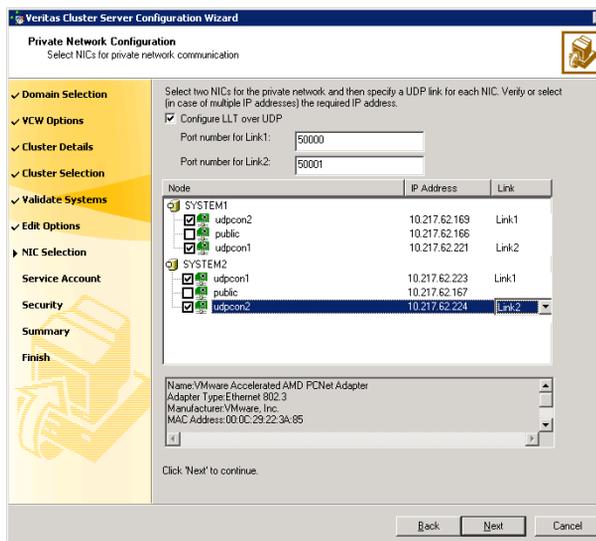
- To configure the VCS private network over Ethernet



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

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

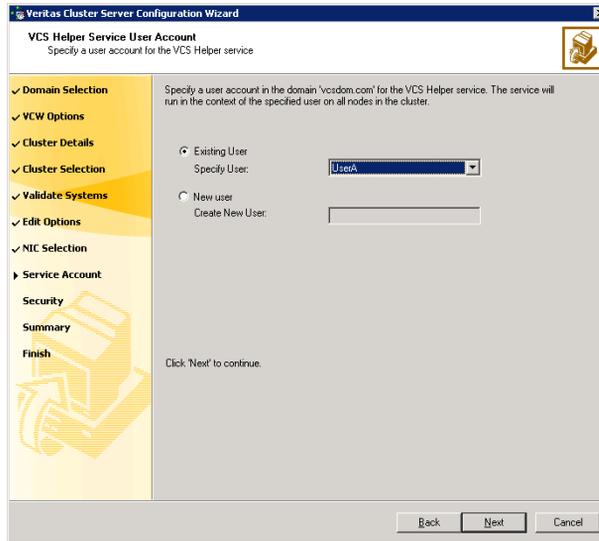
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 802, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

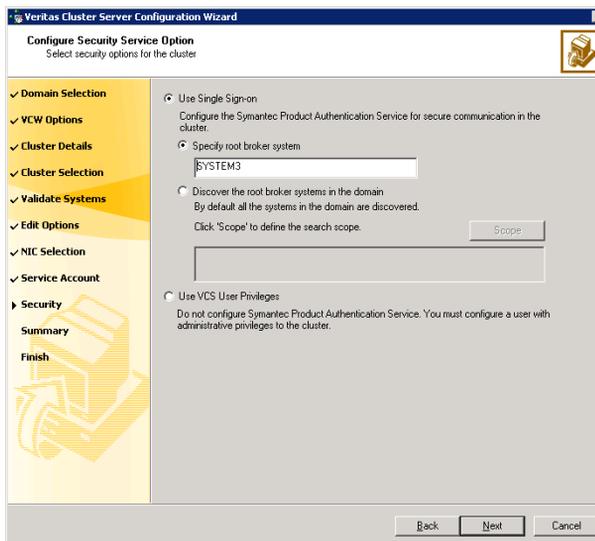
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

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

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 16-4](#) on page 810 contains some more examples of search criteria.

Table 16-4 Search criteria examples

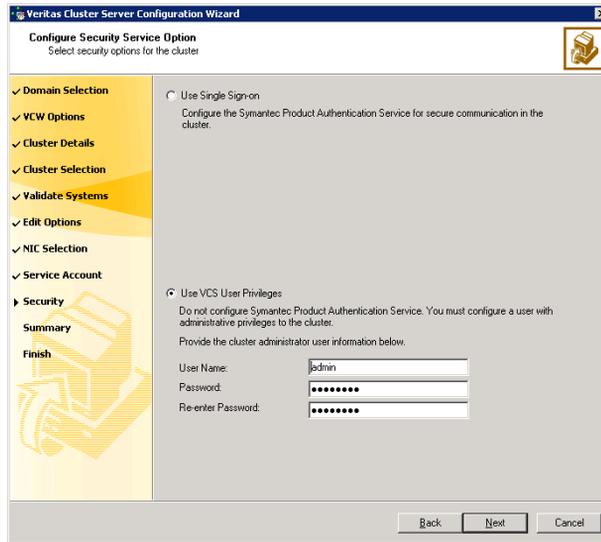
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

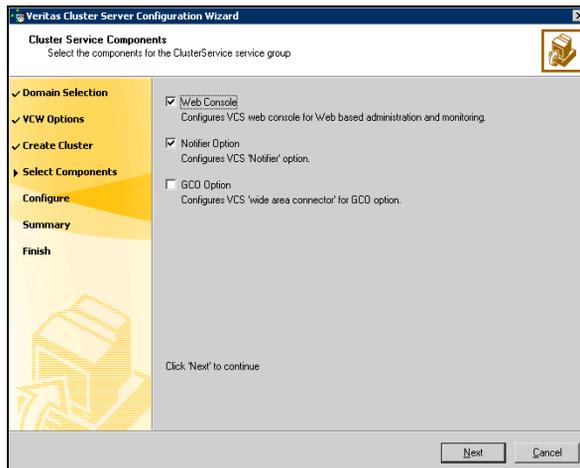
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See “[Configuring Web console](#)” on page 813.

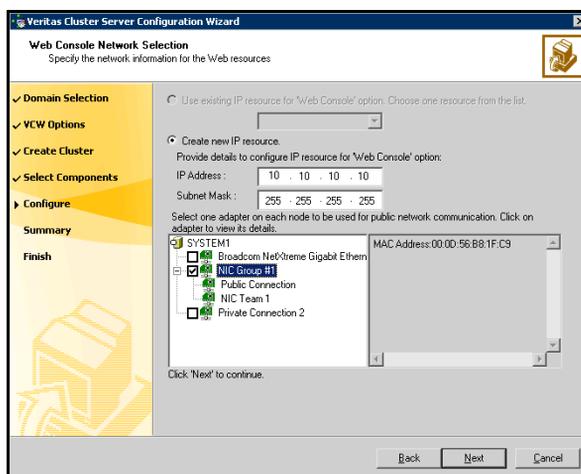
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 814.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



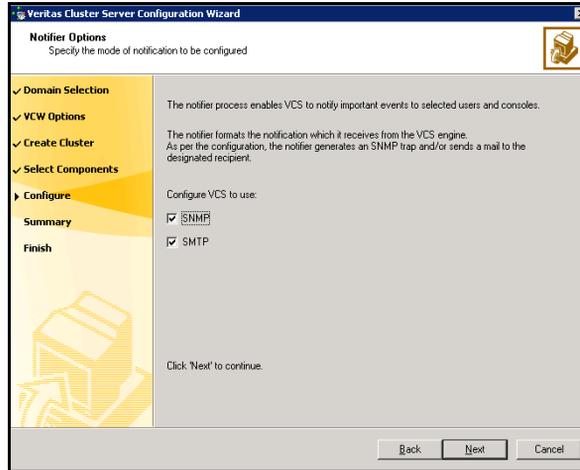
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 814. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

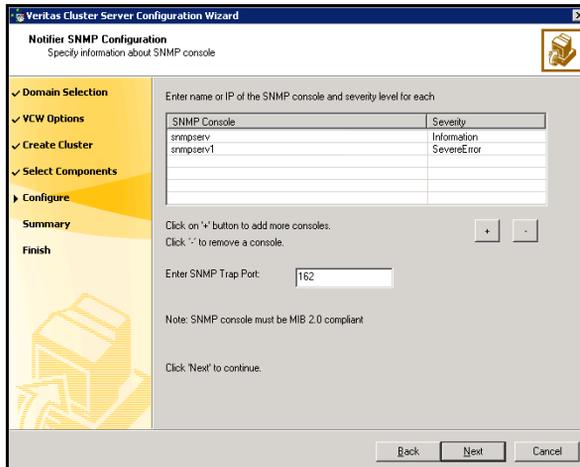
To configure notification

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

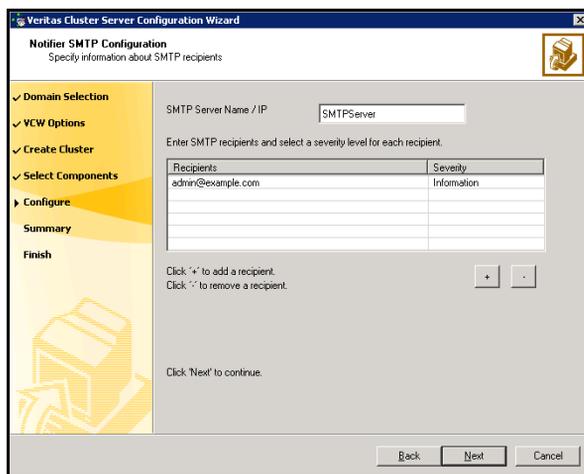


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

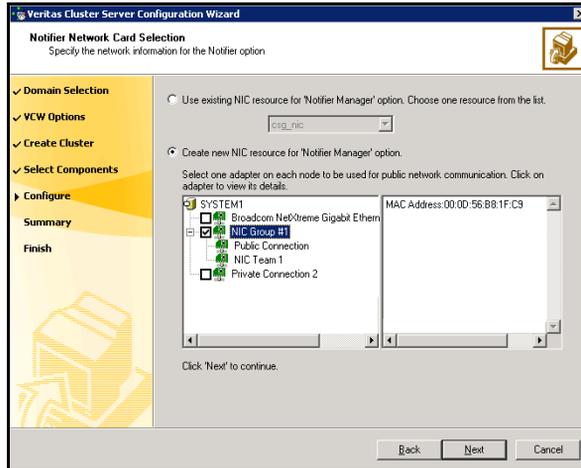


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Verifying your primary site configuration

Before you begin configuring disaster recovery, make sure that SQL Server 2000 has been configured for high availability at the primary site. If you have not yet configured SQL for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [Chapter 4, “Deploying SFW HA for high availability: New SQL Server 2000 installation”](#) on page 49.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center. See [Chapter 13, “Configuring Replicated Data Clusters for SQL 2000”](#) on page 539.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Veritas Volume Replicator (VVR)
- EMC SRDF
- Hitachi TrueCopy

For array-based hardware replication, you can use any replication agent supported by Veritas Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only”](#) on page 870.

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:

- [“Setting up security for VVR”](#) on page 818
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 821
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 823

Setting up security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you configure the VxSAS service on all cluster nodes on both the primary and secondary sites.

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator’s Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxscfg.exe` from the command prompt of the required machine.
Read the information provided on the Welcome page and click **Next**.

2 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 the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts. Click **Next**.

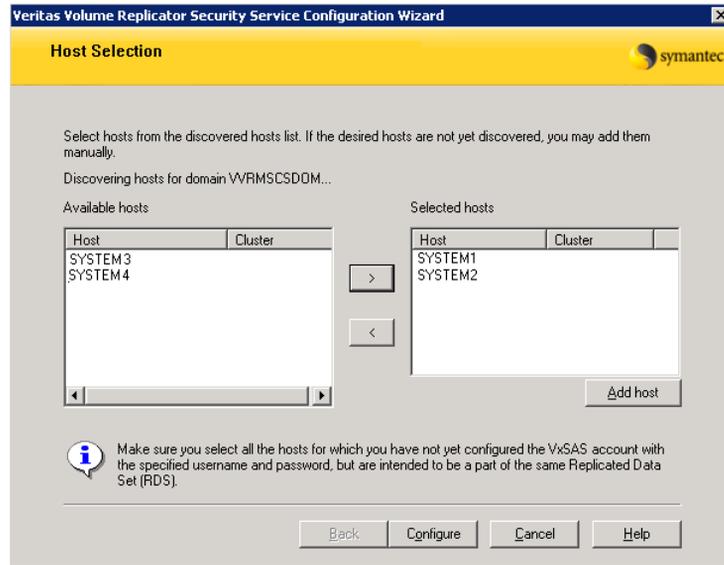
3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- 6 Click **Finish** to exit the wizard.

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for configuring EMC SRDF”](#) on page 821
- [“Replication requirements for EMC SRDF”](#) on page 821

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
- On the secondary site, the wizard verifies that no devices are RDF1.

Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

In addition, note the following agent requirement:

- Device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

Before using the DR wizard, review the following topics:

- “[Software requirements for Hitachi TrueCopy](#)” on page 823
- “[Replication requirements for Hitachi TrueCopy](#)” on page 823

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

For details, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.
- Enter the primary and secondary site file paths for the horcm files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:
`System Driver\Windows`
- The horcm files are named `horcmnn.conf` (where `nn` is a positive number without a leading zero, for example, `horcm1.conf` but not `horcm01.conf`).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

In order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the SQL Server service group as well as any dependent service groups except for the RVG service group.

See the *Veritas Cluster Server Administrator's Guide*.

To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Add the user. Specify the name in the format `username@domain`.

```
hauser -add user [-priv <Administrator|Operator>]
```

- 3 Modify the attribute of the service group to add the user. Specify the SQL Server service group and any dependent service groups except for the RVG service group.

```
hauser -add user [-priv <Administrator|Operator> [-group  
service_groups]]
```

- 4 Reset the configuration to read-only:

```
haconf -dump -makero
```

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Add the user. Specify the name in the format `username@domain`.

```
hauser -add user [-priv <Administrator|Operator>]
```

- 3 Reset the configuration to read-only:

```
haconf -dump -makero
```

Configuring disaster recovery with the DR wizard

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (VVR replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure VVR replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

You will need to exit the wizard after the storage cloning task to install the SQL application. The wizard allows you to exit after the logical completion of each task.

Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking **Next** then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.
- The clusters taking part in the DR configuration should have distinct names.

- After SQL is installed on the secondary site, 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.
- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- One static IP address is available per application service group to be cloned.
- If using VVR for replication, a minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- [“Setting up security for VVR”](#) on page 818
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 821
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 823

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

Note: By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

- 2 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.
- 3 In the System Selection panel, complete the requested information:

System Name Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the SQL instance is online.

 If you have launched the wizard on the system where the instance is online at the primary site, you can also specify `localhost` to connect to the system.

Click **Next**.

- 4 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.
You can choose to clone only the parent service group by not selecting the dependent service group. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported.
The panel lists only service groups that contain a MountV resource.
Click **Next**.
- 5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.
Click **Next**.

- 6 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)	<p>Select this option if you want to configure VVR replication.</p> <p>Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment.</p> <p>The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p> <p>You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.</p>
Configure EMC SRDF and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.</p> <p>Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.</p> <p>Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>

Configure the Global Cluster Option (GCO) only

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 VVR replication environment.

Click **Next**.

- 7 Continue with the next DR configuration task.
For VVR replication, see “[Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)](#)” on page 831.
For array-based replication, see “[Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)](#)” on page 835.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic for the wizard prerequisites before continuing with the storage cloning procedure:

[“Configuring disaster recovery with the DR wizard”](#) on page 826

To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.
- 2 Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.

Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. <ul style="list-style-type: none">■ If the volume does not exist, a new volume will be created.■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.
--------------------	--

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3 In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the

primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.

Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4 In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group	Displays the disk group name to which the volume belongs.
Volume (Volume Size)	Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.
Available Disks	Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> option to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.
	Select disks for each unavailable volume that you want to clone on to the secondary.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.
View Primary Layout	Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.

- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the SQL Server Installation panel, review the information and do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the application on the required nodes. 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. After completing the application installation, you can launch the DR wizard again.
 - Click **Next** to continue with service group cloning if the application is already installed on the required nodes.
 - If the DR wizard is run from a remote node, you can keep the wizard running on that node. You can then install the SQL application locally on each of the required nodes and then click **Next** to continue.
 - If you are running the DR wizard from a local system and need to install the SQL application on that system, the system gets restarted when the application installation is complete. You can then restart the wizard.

If you exit the wizard at any point and then restart it, the wizard starts from the Welcome panel. Continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. The wizard proceeds to the storage cloning panel. If it detects that the storage is identical, it proceeds to the service group cloning.

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 size of the volume at the primary site, whichever is smaller. The wizard deletes the temporary configuration after application installation.

If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning check box on the Storage Cloning panel.

If you have not yet started the wizard, see the following topic for the wizard prerequisites before continuing with the storage cloning procedure:

[“Configuring disaster recovery with the DR wizard”](#) on page 826.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface Default: C:\HORCM\etc where C is the system drive.
-----------------------	--

HORCM files location Path to the horcm configuration files (horcm nn .conf)
Default: C:\Windows
where C is the system drive
An horcm configuration file is required by the RAID Manager on all nodes; however the wizard does not validate this.

- 4 In the Storage Cloning panel, 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**.

- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the SQL Server Installation panel, review the information and do one of the following:
 - 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 are running the DR Wizard from a local system and need to install the SQL application on that system, click **Finish** to exit the wizard and proceed with installing the application on the required nodes.

After completing the application installation, you can launch the DR Wizard again to proceed with service group cloning. At this point 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.
 - If the DR Wizard is run from a remote node, you can keep the wizard running on that node. You can then install the SQL application locally on each of the required nodes.

After completing the application installation, click **Next** to proceed with service group cloning. At this point 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 SQL Server 2000 on the first node (secondary site)

Complete the following procedures to install and configure Microsoft SQL Server 2000:

- [Installing Microsoft SQL Server](#)
- [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the DR Wizard does not mount the volumes on the secondary site and you must format the volumes and mount them manually. See “[Importing the cluster disk group](#)” on page 844 and “[Mounting the volumes](#)” on page 844.

Install Microsoft SQL Server 2000 on the first node for the SQL Server instance using the installation wizard provided with the product. Use the same instance name as was used on the primary site.

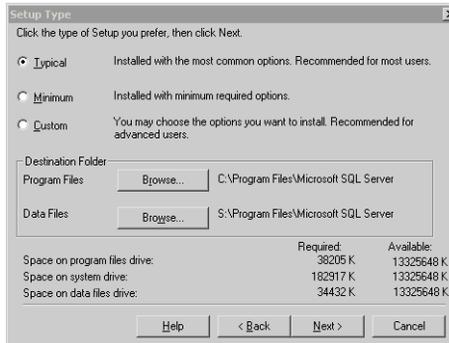
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

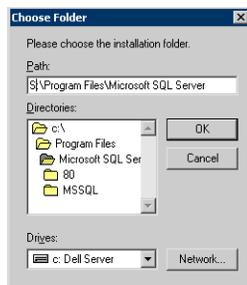
To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.

- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

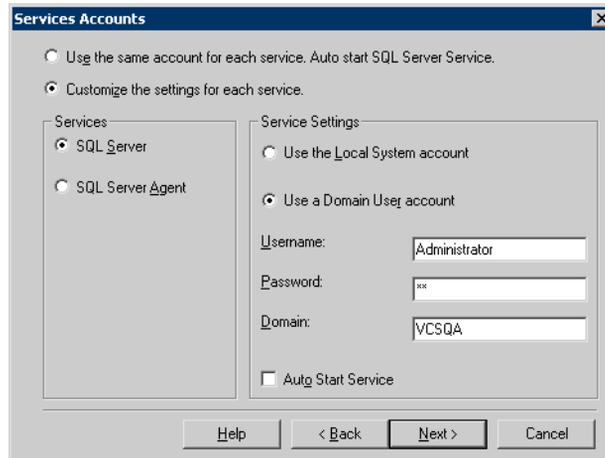


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

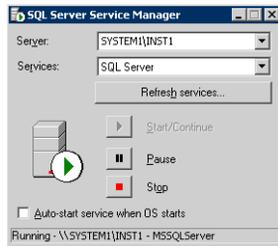
Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node (secondary site)

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 Service”](#) on page 843
- [“Deporting the cluster disk group”](#) on page 843
- [“Importing the cluster disk group”](#) on page 844
- [“Mounting the volumes”](#) on page 844
- [“Renaming shared SQL Server 2000 files”](#) on page 846

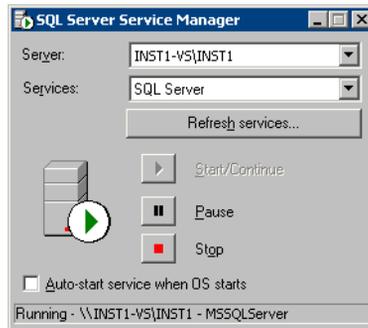
Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 Service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.

- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the Actions menu, select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - Click **OK** in the Import Dynamic Disk Group dialog box.

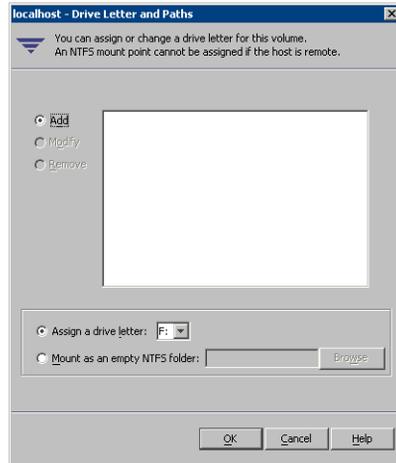
Mounting the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node (secondary site)

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this SQL Server instance:

- [“Installing SQL Server”](#) on page 846
- [“Removing shared SQL Server files”](#) on page 849

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters). See [“Importing the cluster disk group”](#) on page 844 and [“Mounting the volumes”](#) on page 844.

Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

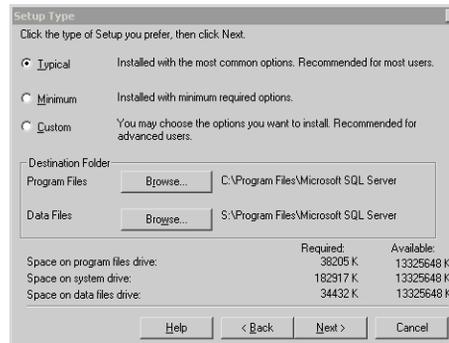
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.

Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.

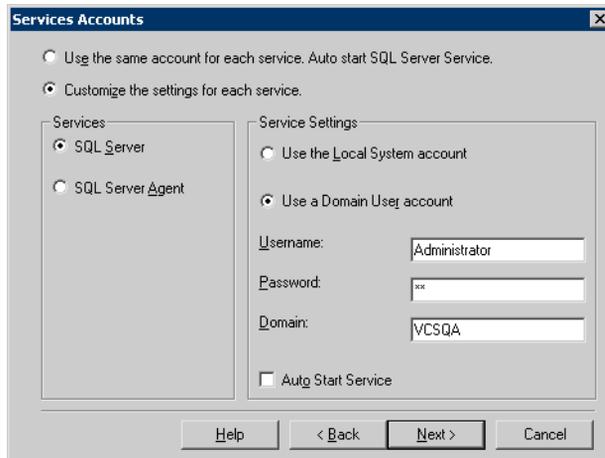
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.

- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node \(secondary site\)](#)” on page 842 and “[Installing SQL Server 2000 on the second node \(secondary site\)](#)” on page 846 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

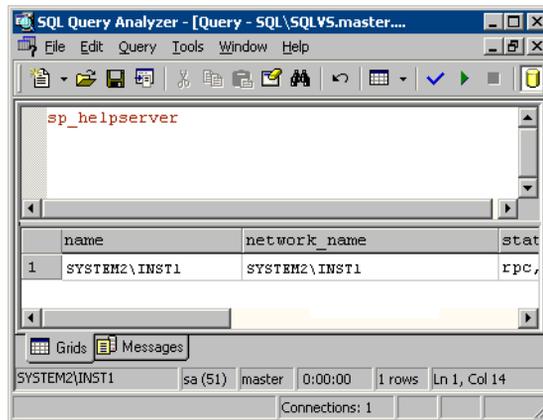
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped checkbox**.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`“sp_dropserver ‘System_Name\Instance_Name.’”`
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
`“sp_dropserver ‘SYSTEM2\INST1.’”`
 For example, for a default instance:
`“sp_dropserver ‘SYSTEM1.’”`
 - Press F5.
- 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter
`"sp_addserver 'Virtual_Server_Name\Instance_Name', local"`
For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.
 - Press F5.

Cloning the service group configuration from the primary to the secondary site

The Disaster Recovery Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Before cloning the service group on the secondary site, verify that you have installed the application on the secondary site on all nodes for this SQL instance.

If you are launching the wizard for the first time, see the following topic for additional information:

["Configuring disaster recovery with the DR wizard"](#) on page 826

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.

If you selected the VVR replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel. If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.

- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:
 - If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.
 - If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.
- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.
- 6 Review the following information displayed in the Service Group Analysis panel and click **Next** to continue with service group cloning.

Service Group Name Displays the list of application-related service groups present on the cluster at the primary site.

Service Group Details on the Primary Cluster Displays the resource attributes for the service group at the primary site. These include:

- IP Resource: consists of the IP address and the subnet mask
- NIC Resource: is the MAC address

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	<p>Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning.</p> <p>Select any additional secondary systems on which you want the wizard to clone the application service group configuration.</p> <p>Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane.</p> <p>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.</p>
Selected Systems	<p>Displays the list of selected systems. The secondary system that you selected earlier in the wizard is listed by default.</p>

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	<p>Displays the list of resources that exist on the primary cluster.</p>
Attribute Name	<p>Displays the attribute name associated with each of the resources displayed in the Resource Name column.</p> <p>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.</p>
Primary Cluster	<p>Displays the primary attribute values for each of the displayed attributes.</p>
Secondary Cluster	<p>The default is the same as the primary cluster. The same virtual IP address can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment. For the MACAddress attribute select the appropriate public NIC from the drop-down list.</p>

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 (✘) 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.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 1171.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication.

Note: By default, in an Exchange or SQL Server environment, the DR wizard organizes all the volumes in a disk group under one Replicated Volume Group (RVG). If you require a different organization, you should configure it using the Veritas Enterprise Administrator (VEA) rather than the DR wizard. For information on setting up VVR replication with the VEA, see [Appendix A, “Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation”](#) on page 1003.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- [“Configuring VVR replication and global clustering”](#) on page 856
- [“Configuring EMC SRDF replication and global clustering”](#) on page 864
- [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 867
- [“Configuring global clustering only”](#) on page 870

Configuring VVR replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure VVR replication and global clustering.

Before you begin, ensure that you have met the following prerequisites:

- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site.

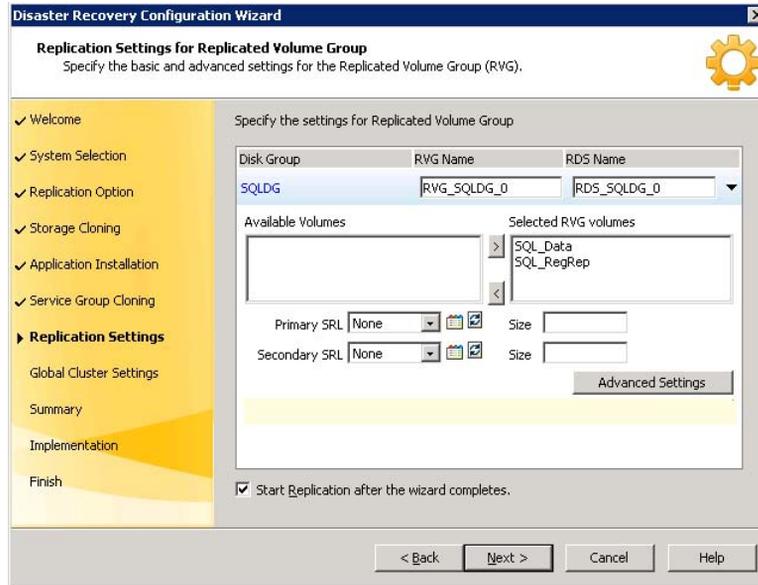
See [“Setting up security for VVR”](#) on page 818

- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure VVR replication and global clustering with the DR wizard.

To configure VVR replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
- 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.



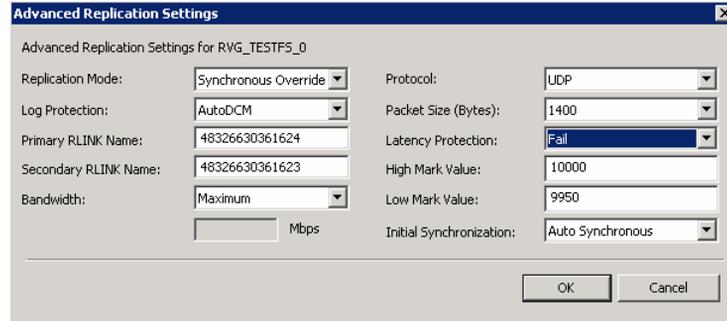
- | | |
|-------------------|---|
| Disk Group | The left column lists the disk groups. By design, an RVG is created for each disk group. |
| RVG Name | Displays the default RVG name. If required, change this to a name of your choice. |
| RDS Name | Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice. |
| Available Volumes | Displays the list of available volumes that have not been selected to be a part of the RVG.

Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane. |

Selected RVG Volumes	<p>Displays the list of volumes that have been selected to be a part of the RVG.</p> <p>To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane.</p> <p>Symantec recommends excluding tempdb from replication. If you earlier moved tempdb to a separate volume in the same disk group as the system database volumes, you can exclude tempdb from replication by removing the tempdb volume from the Selected RVG Volumes pane.</p>
Primary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk.</p> <p>Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Secondary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL.</p> <p>Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Start Replication after the wizard completes	<p>Select this check box to start replication automatically after the wizard completes the necessary configurations.</p> <p>Once replication is configured and running, deselecting the checkbox does not stop replication.</p>

- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator*

Administrator's Guide for additional information on VVR replication options.



Replication Mode Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override**. The default is synchronous override.

Log Protection Select the appropriate log protection from the list. The **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **Off** option disables Replicator Log Overflow protection.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

If the Secondary becomes inactive due to disconnection or administrative action then Replicator log protection is disabled, and the Replicator Log overflows.

The **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, VVR replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit. The default unit is Mega bits 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, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

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**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 5 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

Disk Group	Displays the list of disk groups that have been configured.
RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site.
Copy	Enables you to copy the RLINK attributes across multiple RLINKs. You must have at least two RLINKs to be able to use this operation to copy RLINK attributes from the current to the other RLINKs.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.
Click **Next** to implement the settings.
- 8 In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description

about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have met the prerequisites for replication.

See “[Requirements for EMC SRDF array-based hardware replication](#)” on page 821

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings.

See “[Optional settings for EMC SRDF](#)” on page 866

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.

- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID)	Specify the array ID for the primary site and for the secondary site.
Device Group name	Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.
Available VMDG Resources	Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.

- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.

Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The optional settings use the following defaults:

Option	Default setting
SymHome	C:\Program Files\EMC\SYMCLI\bin
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

Ensure that you have met the prerequisites.

See “[Requirements for Hitachi TrueCopy array-based hardware replication](#)” on page 823

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings.

See “[Optional settings for HTC](#)” on page 870

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 4 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the horcm file is configured properly. If not, you can configure the horcm file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current horcm file.

- 5 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration,

GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

- 10 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only.

For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1 If the wizard is still open after the service group cloning task, continue with the GCO Setup panel. Otherwise, launch the wizard and proceed to the GCO Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 2 In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.
- 3 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.

Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.
-------------------------------	--

- 4 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. Click **Next**.
- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (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.

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- For VVR replication:
 - Ensure VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct

volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.

- Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
- Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
- Confirm that the RVG service groups are online at the primary and secondary sites.
- Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using VVR for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using VVR for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
 - starting a VVR replication checkpoint

- performing a block level backup
- ending the VVR replication checkpoint
- restoring the block level backup at the DR site
- starting replication from the VVR replication checkpoint

To learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.

- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.

- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the -secure option to the StartProgram resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value.
 For example:

```
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe"  
-secure
```
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.
 Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel  
<low|medium|high> [--hashfile <filename> | --hash <root  
hash in hex>]
```

 For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:
 from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low  
from RB2, type:  
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the ClusterService-Proc (wac) resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

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 790.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 16-5 Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent remains online on the primary site. 	1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.
	<ul style="list-style-type: none"> ■ An alert notification at the secondary site occurs for the child service group only. 	2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
	<ul style="list-style-type: none"> ■ The RVG group remains online. 	

Table 16-5 Online, local, soft dependency link

Failure condition	Results	Action required
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Table 16-6 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

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 16-7 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. 	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	<ul style="list-style-type: none"> ■ An alert notification at the secondary site occurs for the child service group only. 	Do not take the RVG group offline at the primary site.
	<ul style="list-style-type: none"> ■ The RVG group remains online. 	
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	<ul style="list-style-type: none"> ■ An alert notification at the secondary site occurs for the parent only. 	<ol style="list-style-type: none"> 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	<ul style="list-style-type: none"> ■ The RVG group remains online. 	

Deploying disaster recovery: New SQL Server 2005 installation

This chapter covers the following topics:

- [Tasks for a new disaster recovery installation of Microsoft SQL Server 2005](#)
- [Reviewing the requirements](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Setting up the secondary site: Configuring SFW HA and setting up a cluster](#)
- [Verifying your primary site configuration](#)
- [Setting up your replication environment](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)](#)
- [Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)](#)
- [Installing and configuring SQL Server 2005 on the first node \(secondary site\)](#)
- [Preparing to install SQL Server 2005 on the second node \(secondary site\)](#)

- [Installing SQL Server 2005 on the second node \(secondary site\)](#)
- [Setting the internal name of the clustered instance](#)
- [Cloning the service group configuration from the primary to the secondary site](#)
- [Configuring replication and global clustering](#)
- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Adding multiple DR sites \(optional\)](#)
- [Recovery procedures for service group dependencies](#)

Tasks for a new disaster recovery installation of Microsoft SQL Server 2005

Before setting up SFW HA disaster recovery at the secondary site, you must complete the high availability configuration on the primary site.

See [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#) on page 199.

You can also configure disaster recovery for a primary site that is configured as a replicated data cluster.

See [Chapter 14, “Configuring Replicated Data Clusters for SQL 2005”](#) on page 655.

After setting up an SFW HA environment for SQL on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

After service group configuration, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [Chapter 2, “Using the Solutions Configuration Center”](#) on page 31.

Note: If you want to configure the secondary site manually, without using the DR wizard, see [“Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation”](#) on page 1075.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171.

[Table 17-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 17-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“Reviewing the requirements” on page 885	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 890	<ul style="list-style-type: none"> ■ Understanding site failover in a DR environment ■ Reviewing the sample configuration ■ Understanding supported disaster recovery configurations for service group dependencies
“Configuring the storage hardware and network” on page 894	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed
“Setting up the secondary site: Configuring SFW HA and setting up a cluster” on page 897	<ul style="list-style-type: none"> ■ Installing SFW HA ■ Configuring the cluster using the Veritas Cluster Server Configuration Wizard
“Verifying your primary site configuration” on page 920	Verifying that SQL has been configured for high availability at the primary site and that the service groups are online
“Setting up your replication environment” on page 920	Ensuring that replication prerequisites for your selected method of replication are met before running the DR wizard
“Assigning user privileges (secure clusters only)” on page 928	For a secure cluster only, assigning user privileges

Table 17-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“Configuring disaster recovery with the DR wizard” on page 929	<ul style="list-style-type: none"> ■ Reviewing prerequisites for the DR wizard ■ Starting the DR wizard and making the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method
“Cloning the storage on the secondary site using the DR wizard (VVR replication option)” on page 933	(VVR replication option) Cloning the storage configuration on the secondary site using the DR wizard
“Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 937	(EMC SRDF, Hitachi TrueCopy, or GCO only replication option) Using the DR wizard to create temporary storage for installation on the secondary site
“Installing and configuring SQL Server 2005 on the first node (secondary site)” on page 941	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2005 ■ Configuring SQL services
“Preparing to install SQL Server 2005 on the second node (secondary site)” on page 946	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters or paths as needed ■ Removing shared SQL files from the cluster disk group
“Installing SQL Server 2005 on the second node (secondary site)” on page 949	Installing SQL Server 2005
“Setting the internal name of the clustered instance” on page 953	Setting the internal name of the clustered instance
“Cloning the service group configuration from the primary to the secondary site” on page 956	Cloning the service group configuration from the primary to the secondary site using the DR wizard

Table 17-1 Tasks for deploying disaster recovery for SQL Server 2005

Objective	Tasks
“Configuring replication and global clustering” on page 960	<ul style="list-style-type: none"> ■ (VVR replication) Using the wizard to configure replication and global clustering ■ (EMC SRDF replication) Setting up replication and then using the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Setting up replication and then using the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Using the wizard to configure global clustering, and then setting up replication
“Verifying the disaster recovery configuration” on page 976	Verifying that the secondary site has been fully configured for disaster recovery
“Establishing secure communication within the global cluster (optional)” on page 978	Adding secure communication between local clusters within the global cluster (optional task)
“Adding multiple DR sites (optional)” on page 980	Optionally, adding additional DR sites to a VVR environment
“Recovery procedures for service group dependencies” on page 981	Bringing the service groups online after failover to the secondary site

Reviewing the requirements

This DR solution requires a primary site and secondary site.

Review the following installation and configuration requirements for your systems before installation. Minimum requirements and Symantec recommended requirements may vary.

- [“Disk space requirements”](#) on page 886
- [“Requirements for Veritas Storage Foundation High Availability for Windows \(SFW HA\)”](#) on page 886

Disk space requirements

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

[Table 17-2](#) on page 886 estimates disk space requirements for SFW HA.

Table 17-2 Disk space requirements

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

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required) ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none"> ■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) ■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none"> ■ Windows Server 2008 for 64-bit Itanium (IA64) ■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 890.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW

HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

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

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Reviewing the configuration

During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Requirements for Veritas Storage Foundation High Availability for Windows \(SFW HA\)”](#) on page 886.

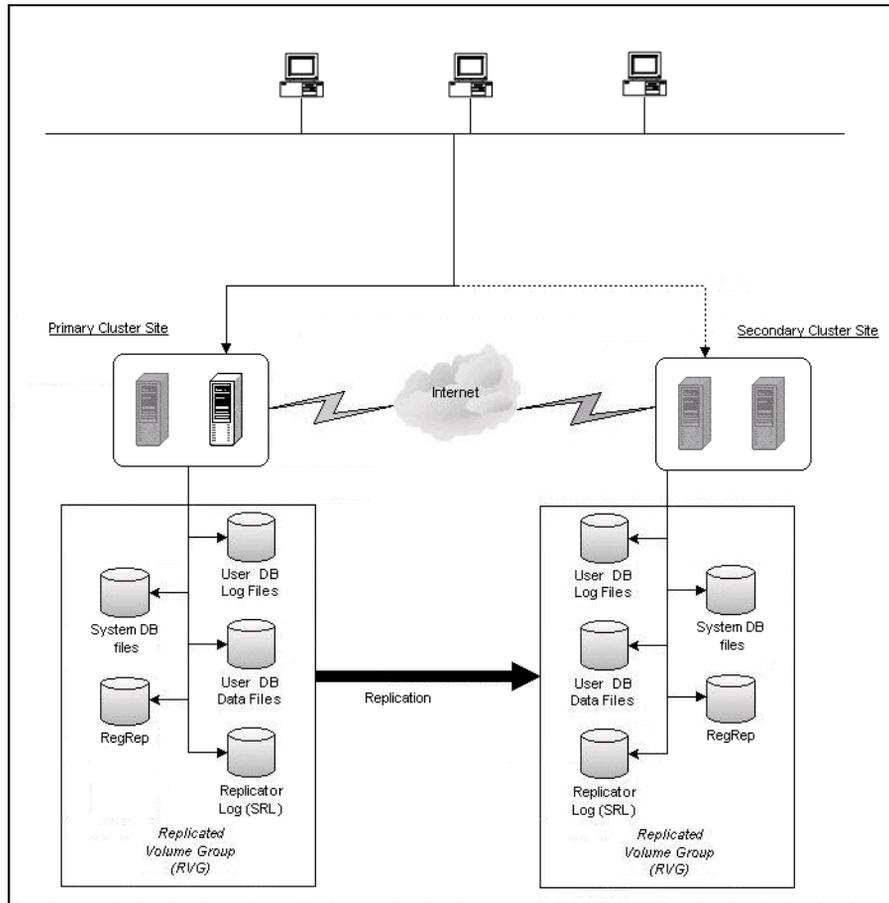
You should have these IP addresses available before you start deploying your environment.

[Figure 17-1](#) illustrates a typical clustered VVR configuration for an active/passive configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG.

You can also configure disaster recovery for an active/active configuration.

See [“Active-Active configuration”](#) on page 212 in [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#).

Figure 17-1 Typical VVR configuration



If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

You can choose to configure replication using VVR or an agent-supported array-based hardware replication. You can use the DR wizard to configure VVR replication or required options for the VCS agents for EMC SRDF or 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.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site. 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:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2005 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	(VVR only) replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components (VVR only)

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

The Disaster Recovery wizard supports only one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

The wizard clones dependent service groups as global groups.

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends

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

- Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
 - 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.

- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.

- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Setting up the secondary site: Configuring SFW HA and setting up a cluster

Install SFW HA and configure the cluster at the secondary site.

Begin with verifying that the requirements are met on the secondary site:

- See [“Reviewing the requirements”](#) on page 885
- Then continue with the following topics:
- [“Installing SFW HA”](#) on page 897
 - [“Configuring the cluster”](#) on page 904

Installing SFW HA

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table 17-3](#) on page 898 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table 17-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 899.

To change the driver signing options on each system

- 1 Log on locally to the system.
 - 2 Open the Control Panel and click **System**.
 - 3 Click the **Hardware** tab and click **Driver Signing**.
 - 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
 - 5 Click **OK**.
 - 6 Repeat for each computer.
- If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

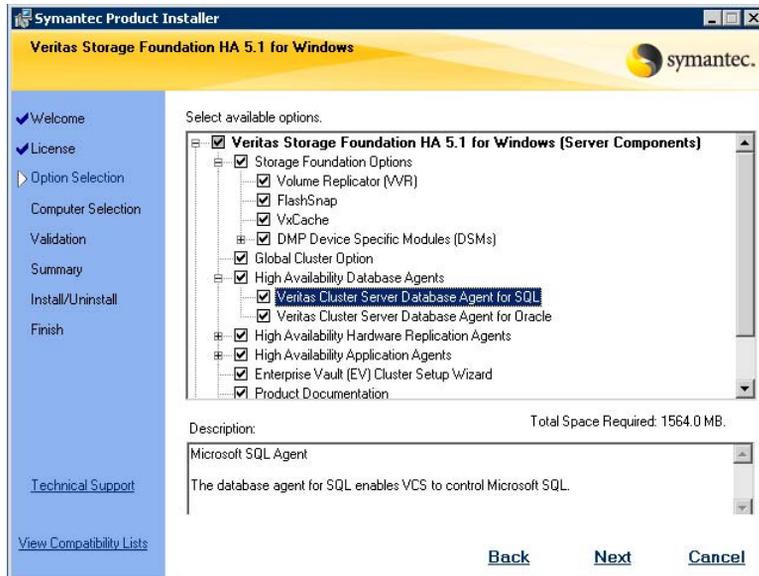
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**

- To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.
- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



Veritas Cluster Server Database Agent for SQL Required to configure high availability for SQL Server.

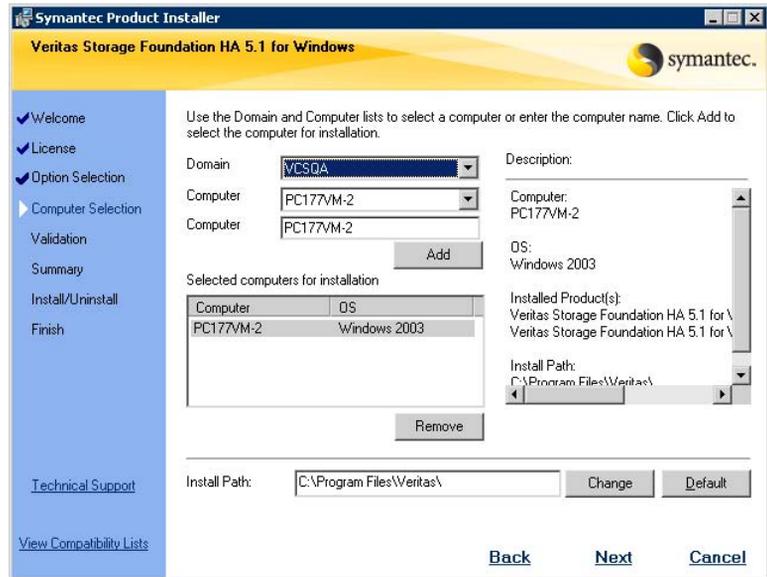
Client Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration.

Global Cluster Option Required for a disaster recovery configuration only.

Veritas Volume Replicator For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR.

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

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

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

Configuring the cluster

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

Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

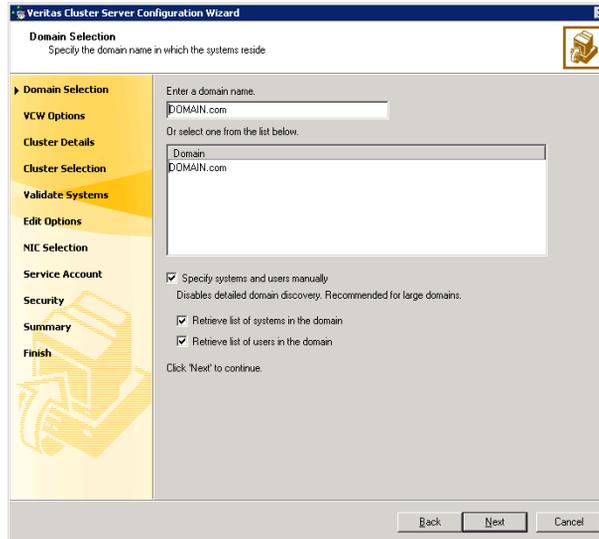
Refer to the *Veritas Cluster Server Administrator's Guide* for complete installation and configuration details on VCS, and additional instructions on removing or modifying cluster configurations.

Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

To configure a VCS cluster

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

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



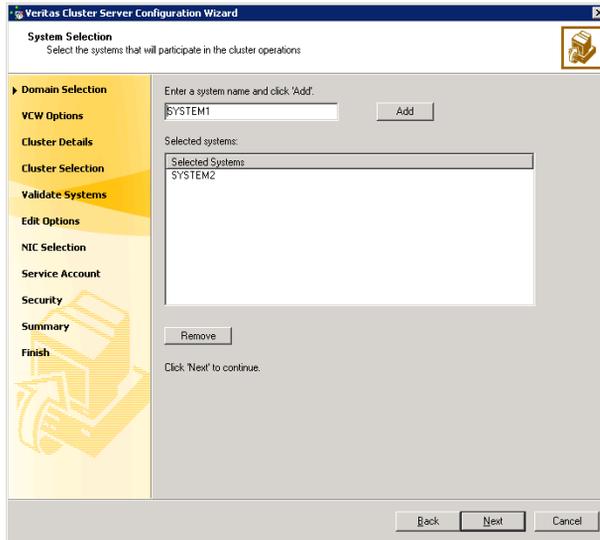
To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
 Proceed to [step 8](#) on page 907.

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

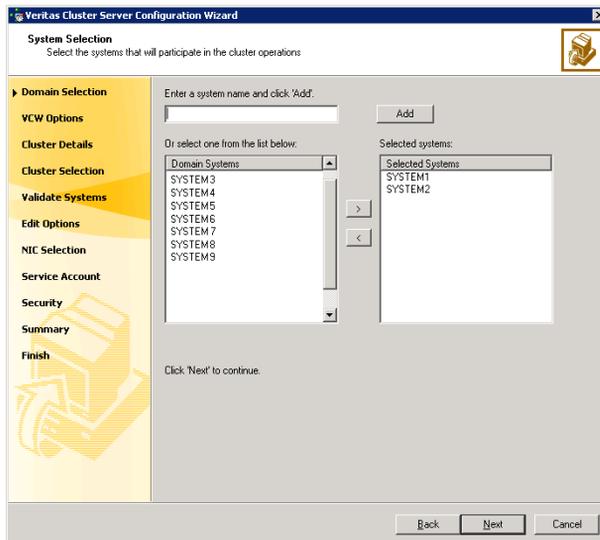
- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
 If you chose to retrieve the list of systems, proceed to [step 6](#) on page 906. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 907.

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



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

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

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

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.

The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window, specifically the 'Cluster Details' step. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Cluster Details - Enter necessary details to create the new cluster'. The left sidebar contains a navigation pane with the following items: Domain Selection, VCS Options, Cluster Details (selected), Cluster Selection, Validate Systems, Edit Options, NIC Selection, Service Account, Security, Summary, and Finish. The main area of the wizard is titled 'Cluster Details' and contains the following fields and options:

- Cluster Name: MYCLUSTER
- Cluster ID: 2
- Operating System: Windows 2003 (x86)
- Select all systems:
- Available Systems list: SYSTEM1, SYSTEM2 (both checked)
- Total number of systems selected to create the cluster : 2
- Click 'Next' to continue.

At the bottom of the window, there are three buttons: Back, Next, and Cancel.

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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 905 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

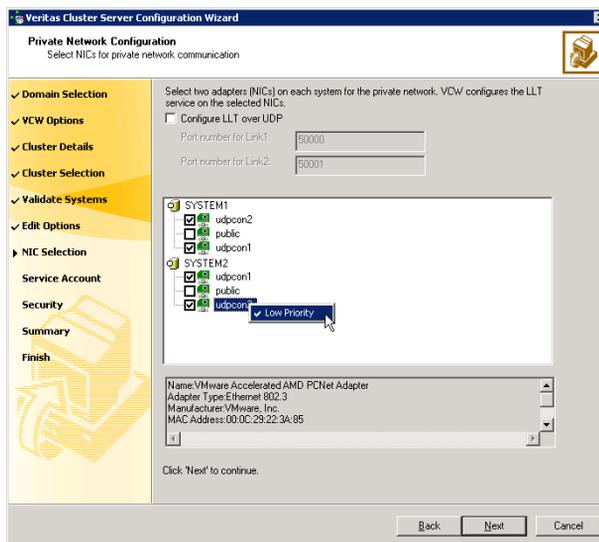
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 907, proceed to the next step. Otherwise, proceed to [step 12](#) on page 911.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

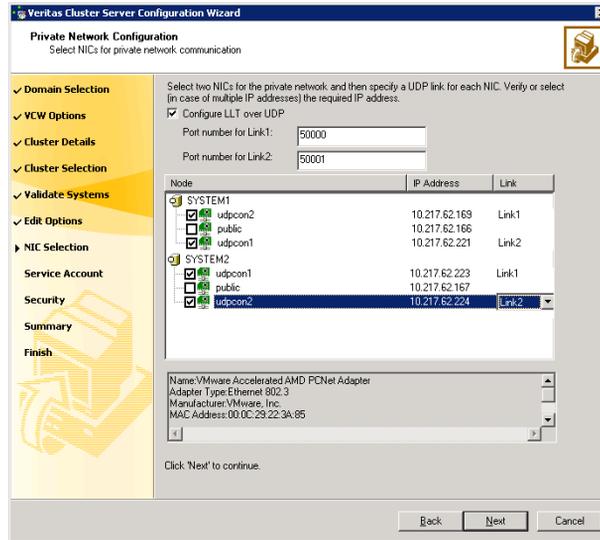
- To configure the VCS private network over Ethernet



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

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

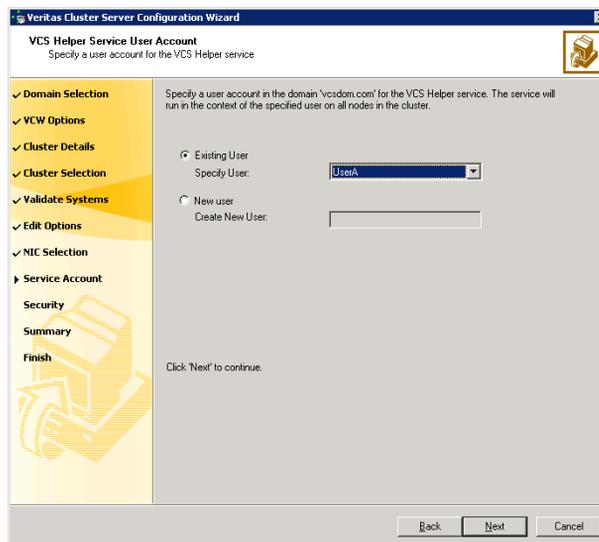
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list,
 - If you chose not to retrieve the list of users in [step 4](#) on page 905, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

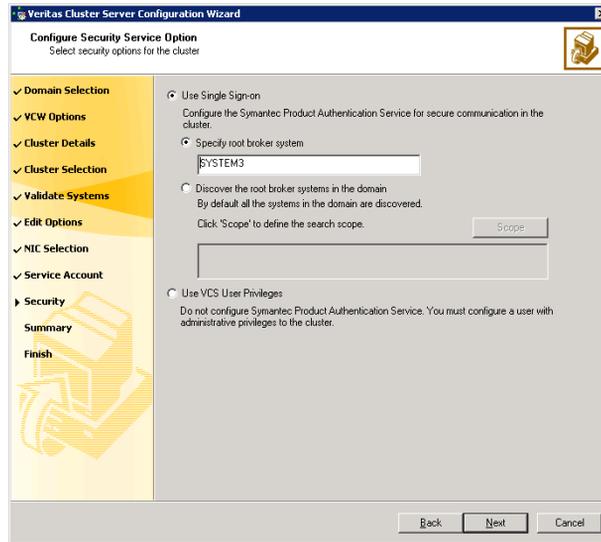
Do not append the domain name to the user name; do not type the user name as DOMAIN\user or user@DOMAIN.

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

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

Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities.

These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. [Table 17-4](#) on page 913 contains some more examples of search criteria.

Table 17-4 Search criteria examples

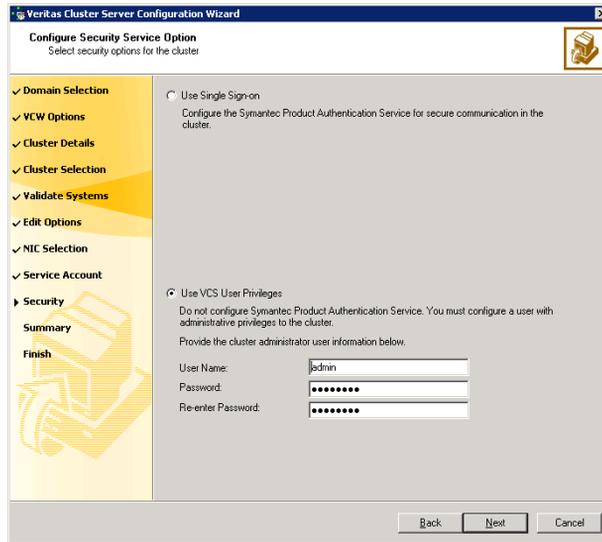
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege:



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

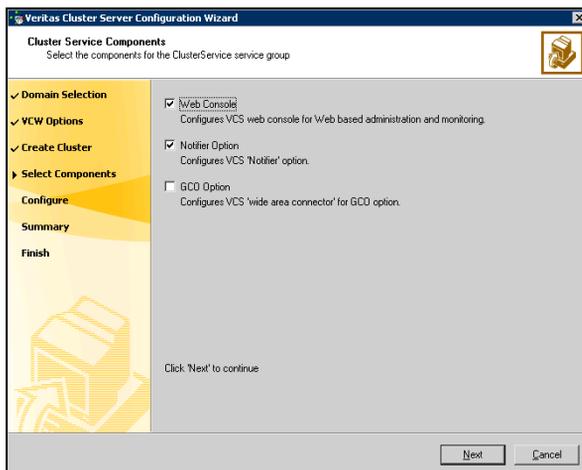
At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and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

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

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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



- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 916.

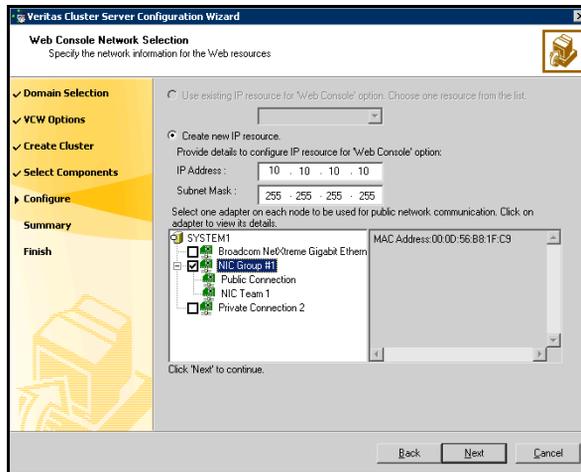
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See “[Configuring notification](#)” on page 917.

Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



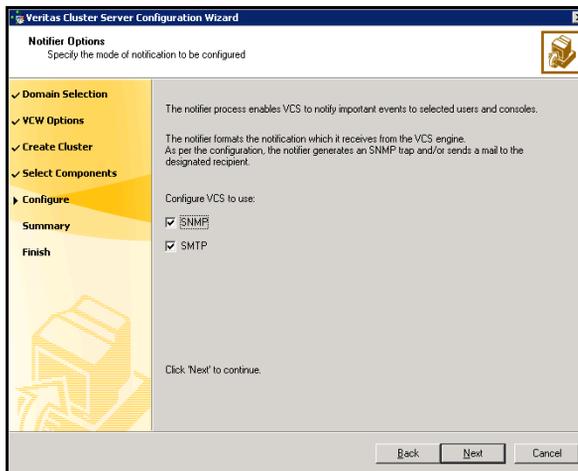
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to: “[Configuring notification](#)” on page 917. Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

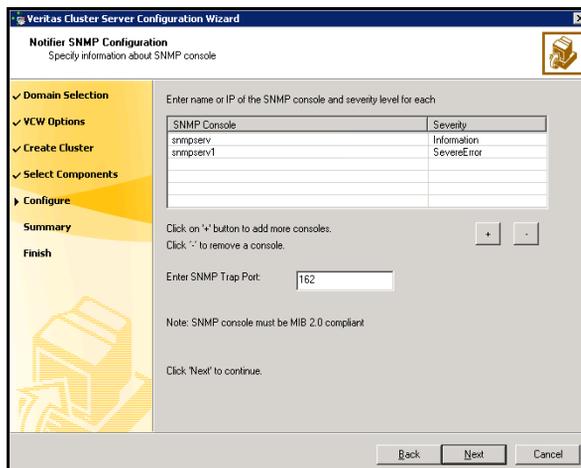
To configure notification

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

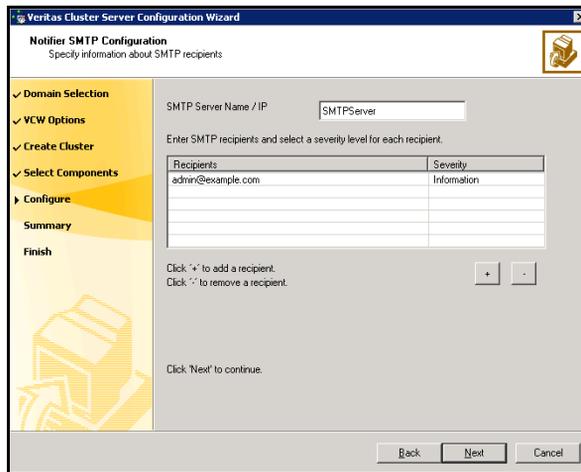


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

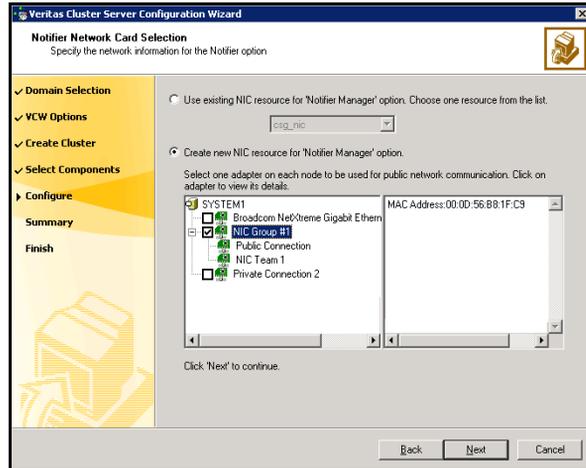


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Verifying your primary site configuration

Before you begin configuring disaster recovery, make sure that SQL Server 2005 has been configured for high availability at the primary site. If you have not yet configured SQL for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [Chapter 6, “Deploying SFW HA for high availability: New SQL Server 2005 installation”](#) on page 199.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center. See [Chapter 14, “Configuring Replicated Data Clusters for SQL 2005”](#) on page 655.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Veritas Volume Replicator (VVR)
- EMC SRDF
- Hitachi TrueCopy

For array-based hardware replication, you can use any replication agent supported by Veritas Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only”](#) on page 974.

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:

- [“Setting up security for VVR”](#) on page 921
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 924
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 926

Setting up security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you configure the VxSAS service on all cluster nodes on both the primary and secondary sites.

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxscfg.exe` from the command prompt of the required machine.

Read the information provided on the Welcome page and click **Next**.

2 Complete the Account Information panel as follows:

Account name Enter the administrative account name.
(domain\account)

Password Specify a password.

If you have already configured the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

Click **Next**.

3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

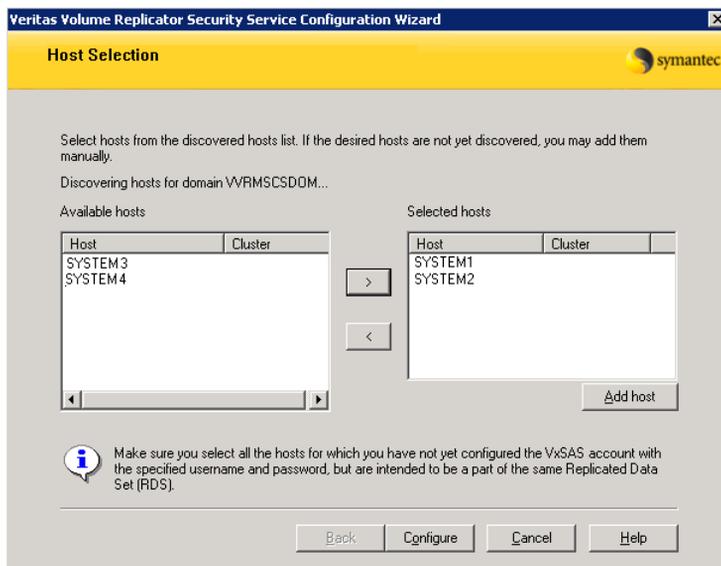
Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

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

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for configuring EMC SRDF”](#) on page 924
- [“Replication requirements for EMC SRDF”](#) on page 924

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
- On the secondary site, the wizard verifies that no devices are RDF1.

Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.
- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

In addition, note the following agent requirement:

- Device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

Before using the DR wizard, review the following topics:

- [“Software requirements for Hitachi TrueCopy”](#) on page 926
- [“Replication requirements for Hitachi TrueCopy”](#) on page 926

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

For details, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.
- Enter the primary and secondary site file paths for the horcm files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:
`System Driver\Windows`
- The horcm files are named `horcmnn.conf` (where *nn* is a positive number without a leading zero, for example, `horcm1.conf` but not `horcm01.conf`).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

In order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the SQL Server service group as well as any dependent service groups except for the RVG service group.

See the *Veritas Cluster Server Administrator's Guide*.

To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Modify the attribute of the service group to add the user. Specify the SQL Server service group and any dependent service groups except for the RVG service group.
`hauser -add user [-priv <Administrator|Operator> [-group service_groups]]`
- 4 Reset the configuration to read-only:
`haconf -dump -makero`

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Reset the configuration to read-only:
`haconf -dump -makero`

Configuring disaster recovery with the DR wizard

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (VVR replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure VVR replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

You will need to exit the wizard after the storage cloning task to install the SQL application. The wizard allows you to exit after the logical completion of each task.

Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking **Next** then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.
- The clusters taking part in the DR configuration should have distinct names.

- After SQL is installed on the secondary site, 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.
- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- One static IP address is available per application service group to be cloned.
- If using VVR for replication, a minimum of one static IP address per site is available for each application instance running in the cluster.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- [“Setting up security for VVR”](#) on page 921
- [“Requirements for EMC SRDF array-based hardware replication”](#) on page 924
- [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 926

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

Note: By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

2 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.

3 In the System Selection panel, complete the requested information:

System Name Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the SQL instance is online.

If you have launched the wizard on the system where the instance is online at the primary site, you can also specify `localhost` to connect to the system.

Click **Next**.

4 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.

You can choose to clone only the parent service group by not selecting the dependent service group. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.

Click **Next**.

- 6 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)	<p>Select this option if you want to configure VVR replication.</p> <p>Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment.</p> <p>The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p> <p>You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.</p>
Configure EMC SRDF and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.</p> <p>Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.</p> <p>Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>

Configure the Global Cluster Option (GCO) only 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 VVR replication environment.

Click **Next**.

- 7 Continue with the next DR configuration task.
For VVR replication, see [“Cloning the storage on the secondary site using the DR wizard \(VVR replication option\)”](#) on page 933.
For array-based replication, see [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 937.

Cloning the storage on the secondary site using the DR wizard (VVR replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic for the wizard prerequisites before continuing with the storage cloning procedure:

[“Configuring disaster recovery with the DR wizard”](#) on page 929.

To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.

- 2 Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.
Recommended Action	Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary. <ul style="list-style-type: none">■ If the volume does not exist, a new volume will be created.■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you

can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3 In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks	For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.
	Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4 In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group	Displays the disk group name to which the volume belongs.
Volume (Volume Size)	Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.
Available Disks	Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> option to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group. Select disks for each unavailable volume that you want to clone on to the secondary.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.

View Primary Layout Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.
- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the SQL Server Installation panel, review the information and do one of the following:
 - Click **Finish** to exit the wizard and proceed with installing the application on the required nodes. 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. After completing the application installation, you can launch the DR wizard again.
 - Click **Next** to continue with service group cloning if the application is already installed on the required nodes.
 - If the DR wizard is run from a remote node, you can keep the wizard running on that node. You can then install the SQL application locally on each of the required nodes and then click **Next** to continue.
 - If you are running the DR wizard from a local system and need to install the SQL application on that system, the system gets restarted when the application installation is complete. You can then restart the wizard.

If you exit the wizard at any point and then restart it, the wizard starts from the Welcome panel. Continue through the wizard, specifying the primary

site system, the service group, the secondary site system, and the replication method. The wizard proceeds to the storage cloning panel. If it detects that the storage is identical, it proceeds to the service group cloning.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

To enable you to install applications, the DR wizard can create a temporary disk group, DR_APP_INSTALL_DG, which contains the volumes and mount points for use in application installation. The temporary configuration uses 500 MB volumes or the volume size at the primary site, depending on which is smaller. The wizard deletes the temporary configuration after application installation. If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning box on the Storage Cloning panel.

If you have not yet started the wizard, see the following topic for the wizard prerequisites before continuing with the procedure:

[“Configuring disaster recovery with the DR wizard”](#) on page 929.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface Default: C:\HORCM\etc where C is the system drive.
-----------------------	--

HORCM files location Path to the horcm configuration files (horcm nn .conf)
Default: C:\Windows
where C is the system drive
An horcm configuration file is required by the RAID Manager on all nodes; however the wizard does not validate this.

- 4 In the Storage Cloning panel, 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**.

- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the SQL Server Installation panel, review the information and do one of the following:
 - 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 are running the DR Wizard from a local system and need to install the SQL application on that system, click **Finish** to exit the wizard and proceed with installing the application on the required nodes.

After completing the application installation, you can launch the DR Wizard again to proceed with service group cloning. At this point 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.
 - If the DR Wizard is run from a remote node, you can keep the wizard running on that node. You can then install the SQL application locally on each of the required nodes.

After completing the application installation, click **Next** to proceed with service group cloning. At this point 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 SQL Server 2005 on the first node (secondary site)

Complete the following procedures to install and configure Microsoft SQL Server 2005:

- [Installing SQL Server 2005 on the first node](#)
- [Setting SQL Server 2005 services to manual start](#)

Installing SQL Server 2005 on the first node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the first node and the volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the DR Wizard does not mount the volumes on the secondary site and you must format the volumes and mount them manually. See “[Importing the cluster disk group](#)” on page 947 and “[Mounting the volumes](#)” on page 948.

Install Microsoft SQL Server 2005 on the first node for the SQL Server instance using the installation wizard provided with the product. Use the same instance name as on the primary site.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005

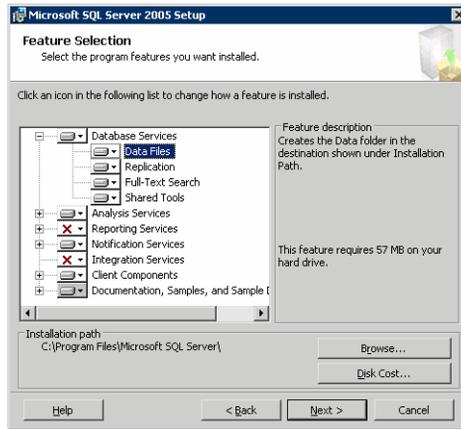
- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.

- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.

If you install optional components on one node, install the same components in the same order on other nodes.

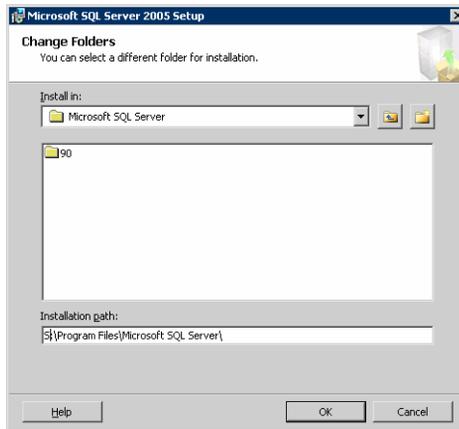
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.



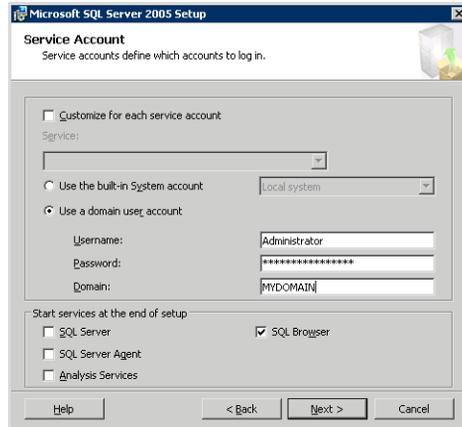
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 942, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group. See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Setting SQL Server 2005 services to manual start

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node (secondary site)

Follow the procedures provided in this section before installing SQL Server on additional nodes for this instance:

- [“Stopping the SQL Server 2005 Service”](#) on page 946
- [“Deporting the cluster disk group”](#) on page 947
- [“Importing the cluster disk group”](#) on page 947
- [“Mounting the volumes”](#) on page 948
- [“Renaming shared SQL Server 2005 files”](#) on page 949

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM3) and then import it to the desired node (SYSTEM4).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM3).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM4).

To import a cluster disk group

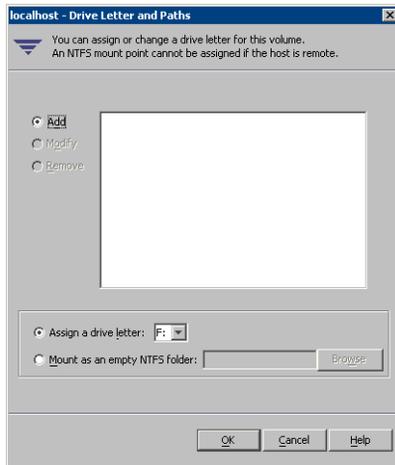
- 1 From the **Actions** menu select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - a In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - b Click **OK** in the **Import Dynamic Disk Group** dialog box.

Mounting the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node (secondary site)

Follow the procedures provided in this section to install and configure SQL Server on additional nodes for this SQL instance:

- [“Installing SQL Server on the second node”](#) on page 949
- [“Removing shared SQL Server files”](#) on page 953

Installing SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters). See [“Importing the cluster disk group”](#) on page 947 and [“Mounting the volumes”](#) on page 948.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

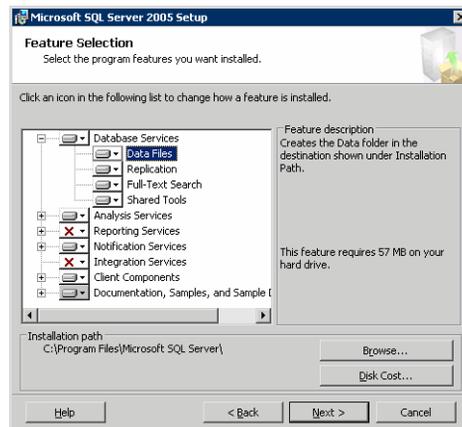
Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

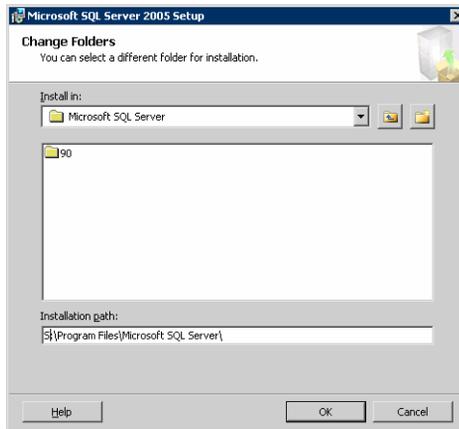
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.



- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 950, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.

See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.
- Refer to “[Setting SQL Server 2005 services to manual start](#)” on page 945 for the procedure.

Repeat the procedures described in “[Preparing to install SQL Server 2005 on the second node \(secondary site\)](#)” on page 946 and “[Installing SQL Server 2005 on the second node \(secondary site\)](#)” on page 949 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

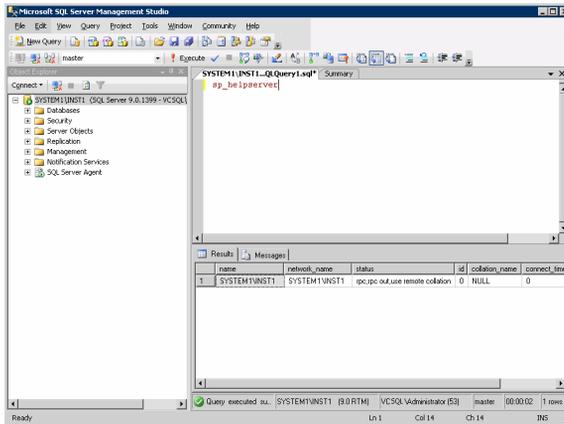
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1 \ INST1. (For a default instance, you see only *System_Name* .)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
`sp_dropserver "System_Name\Instance_Name"`
 where *System_Name\Instance_Name* is the name noted in [step 3](#) on page 954.
 For example, for a named instance:
`sp_dropserver "SYSTEM1\INST1"`
 For example, for a default instance:
`sp_dropserver "SYSTEM1"`
 - Press **F5**.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.
- 9 Stop the SQL instance on the node.
See "[Stopping the SQL Server 2005 Service](#)" on page 946.

Cloning the service group configuration from the primary to the secondary site

The Disaster Recovery Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes for this SQL instance within the cluster, simultaneously.

Before cloning the service group on the secondary site, verify that you have installed the application on the secondary site on all nodes for this SQL instance.

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.

If you are launching the wizard for the first time, see the following topic for additional information:

["Configuring disaster recovery with the DR wizard"](#) on page 929.

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.

- 2 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.
 If you selected the VVR replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel.
 If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.
- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:
 - If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.
 - If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.
- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.
- 6 Review the following information displayed in the Service Group Analysis panel and click **Next** to continue with service group cloning.

Service Group Name Displays the list of application-related service groups present on the cluster at the primary site.

Service Group Details on the Primary Cluster Displays the resource attributes for the service group at the primary site. These include:

- IP Resource: consists of the IP address and the subnet mask
- NIC Resource: is the MAC address

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	<p>Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning.</p> <p>Select any additional secondary systems on which you want the wizard to clone the application service group configuration.</p> <p>Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane.</p> <p>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.</p>
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	<p>Displays the attribute name associated with each of the resources displayed in the Resource Name column.</p> <p>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.</p>
Primary Cluster	Displays the primary attribute values for each of the displayed attributes.

Secondary Cluster The default is the same as the primary cluster. The same virtual IP address can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment. For the MACAddress attribute select the appropriate public NIC from the drop-down list.

Click **Next**.

- 9 In the Service Group Summary, review the attribute information that will be cloned on to the secondary cluster. Click **Back** to change any of the secondary service group attributes. Otherwise, click **Next** to proceed with cloning the service group configuration on the secondary site.
- 10 In the Implementation panel, wait until all the tasks are completed. The progress bar indicates the status of the tasks. Successful tasks are marked with a check symbol. If some task could not be completed successfully, the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**
- 11 If the cloning failed, review the troubleshooting information. Otherwise, click **Next** to continue with the replication and GCO configuration, or with GCO only, depending on which option you selected.
Optionally, you can exit the wizard at this point and launch the wizard again later. When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. Click **Next** to continue to the replication and/or GCO configuration task.
To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using VVR or an agent-supported array-based hardware replication.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- [“Configuring VVR replication and global clustering”](#) on page 960
- [“Configuring EMC SRDF replication and global clustering”](#) on page 968
- [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 971
- [“Configuring global clustering only”](#) on page 974

Configuring VVR replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure VVR replication and global clustering.

Note: By default, in an Exchange or SQL Server environment, the DR wizard organizes all the volumes in a disk group under one Replicated Volume Group (RVG). If you require a different organization, you should configure it using the Veritas Enterprise Administrator (VEA) rather than the DR wizard. For information on setting up VVR replication with the VEA, see [Appendix B, “Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation”](#) on page 1075.

Before you begin, ensure that you have met the following prerequisites:

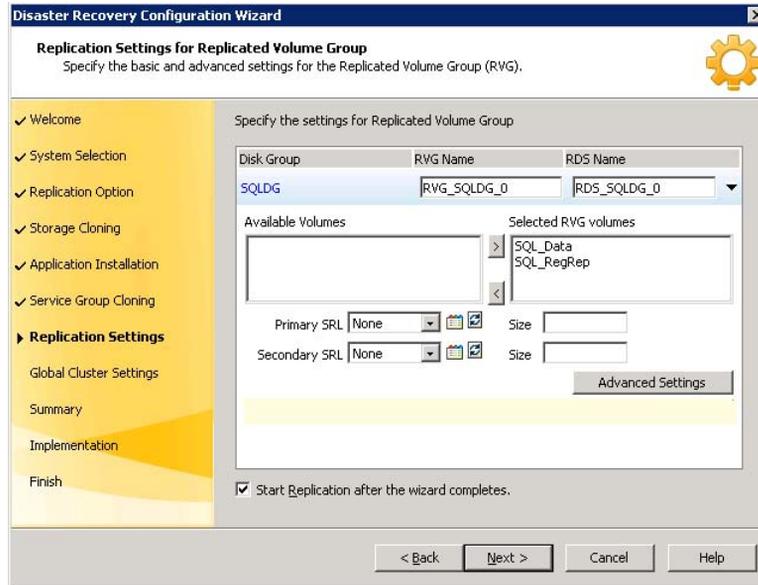
- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site. See the following topic:
 - [“Setting up security for VVR”](#) on page 921

- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure VVR replication and global clustering with the DR wizard.

To configure VVR replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
- 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.

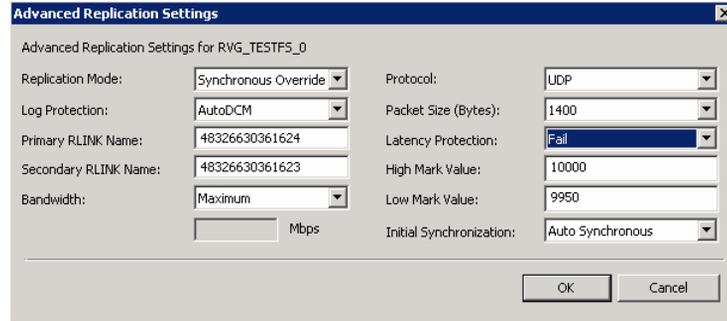


- Disk Group** The left column lists the disk groups. By design, an RVG is created for each disk group.
- RVG Name** Displays the default RVG name. If required, change this to a name of your choice.
- RDS Name** Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice.
- Available Volumes** Displays the list of available volumes that have not been selected to be a part of the RVG.
Either double-click on the volume name or use the > option to move the volumes into the Selected RVG Volumes pane.

Selected RVG Volumes	<p>Displays the list of volumes that have been selected to be a part of the RVG.</p> <p>To remove a selected volume, either double-click the volume name or use the < option to move the volumes into the Available Volumes pane.</p> <p>Symantec recommends excluding tempdb from replication. If you earlier moved tempdb to a separate volume in the same disk group as the system database volumes, you can exclude tempdb from replication by removing the tempdb volume from the Selected RVG Volumes pane.</p>
Primary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk.</p> <p>Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Secondary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL.</p> <p>Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Start Replication after the wizard completes	<p>Select this check box to start replication automatically after the wizard completes the necessary configurations.</p> <p>Once replication is configured and running, deselecting the checkbox does not stop replication.</p>

- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator*

Administrator's Guide for additional information on VVR replication options.



Replication Mode Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override**. The default is synchronous override.

Log Protection Select the appropriate log protection from the list. The **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.

The **Off** option disables Replicator Log Overflow protection.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

If the Secondary becomes inactive due to disconnection or administrative action then Replicator log protection is disabled, and the Replicator Log overflows.

The **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	<p>By default, VVR replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit.</p> <p>The default unit is Mega bits per second (Mbps) and the minimum allowed value is 1 Mbps.</p>
Protocol	Choose TCP or UDP. UDP/IP is the default replication protocol.
Packet Size (Bytes)	Default is 1400 Bytes. From the drop-down list, choose the required packet size for data transfer. The default unit for the packet size is Bytes. You can set the packet size only if the protocol is UDP/IP.
Latency Protection	<p>By default, latency protection is set to Off.</p> <p>When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection.</p> <p>This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.</p>
High Mark Value	<p>This option is enabled only when Latency Protection is set to Override or Fail. It specifies the maximum number of pending updates by which the secondary site can be behind the primary site. The default value is 10000.</p> <p>To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.</p>
Low Mark Value	<p>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.</p>

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, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

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**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 5 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

Disk Group	Displays the list of disk groups that have been configured.
RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC from the drop-down list for the system at the primary and secondary site.
Copy	Enables you to copy the RLINK attributes across multiple RLINKs. You must have at least two RLINKs to be able to use this operation to copy RLINK attributes from the current to the other RLINKs.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.
Click **Next** to implement the settings.
- 8 In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description

about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have met the prerequisites for replication.

See “[Requirements for EMC SRDF array-based hardware replication](#)” on page 924

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings.

See “[Optional settings for EMC SRDF](#)” on page 970

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.

- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID)	Specify the array ID for the primary site and for the secondary site.
Device Group name	Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.
Available VMDG Resources	Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.

- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.

Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for EMC SRDF, Configuration Guide*. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The optional settings use the following defaults:

Option	Default setting
SymHome	C:\Program Files\EMC\SYMCLI\bin
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

Ensure that you have met the prerequisites.

See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 926

In addition, ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings.

See [“Optional settings for HTC”](#) on page 974

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel. Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft

SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 4 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the horcm file is configured properly. If not, you can configure the horcm file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current horcm file.

- 5 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.
- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration,

GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 7 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings. Click **Next**.
- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (✗) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

- 10 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy, Configuration Guide*.

The optional settings use the following defaults:

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only.

For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that the Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1 If the wizard is still open after the service group cloning task, continue with the GCO Setup panel. Otherwise, launch the wizard and proceed to the GCO Setup panel as follows:
 - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
 - In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.
- 2 In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.
- 3 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.

Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.
-------------------------------	--

- 4 In the Settings Summary panel, review the displayed information. Click **Back** if you want to change any of the parameters specified. Click **Next**.
- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check (✓) symbol indicates successful completion of a task. An (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.

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For VVR replication, confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- For VVR replication:
 - Ensure VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct

volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.

- Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
- Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
- Confirm that the RVG service groups are online at the primary and secondary sites.
- Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.
- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Veritas™ Cluster Server Hardware Replication Agent for Hitachi TrueCopy Installation and Configuration Guide*.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using VVR for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using VVR for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
 - starting a VVR replication checkpoint

- performing a block level backup
- ending the VVR replication checkpoint
- restoring the block level backup at the DR site
- starting replication from the VVR replication checkpoint

To learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.

- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover for VVR-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.

- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the -secure option to the StartProgram resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value.
 For example:
`"C:\Program Files\Veritas\Cluster Server\bin\wac.exe"
 -secure`
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.
 Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:
`vssat setuptrust --broker <host:port> --securitylevel
 <low|medium|high> [--hashfile <filename> | --hash <root
 hash in hex>]`
 For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:
 from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low  
from RB2, type:  
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the **ClusterService-Proc (wac)** resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

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 894.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 17-5 Online, local, soft dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent remains online on the primary site. 	1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.
	<ul style="list-style-type: none"> ■ An alert notification at the secondary site occurs for the child service group only. 	2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
	<ul style="list-style-type: none"> ■ The RVG group remains online. 	

Table 17-5 Online, local, soft dependency link

Failure condition	Results	Action required
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Table 17-6 Online, local, firm dependency link

Failure condition	Results	Action required
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ol style="list-style-type: none"> 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

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 17-7 Online, local, hard dependency link

Failure condition	Results	Action required
The child service group fails	■ The parent goes offline on the primary site.	Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ An alert notification at the secondary site occurs for the child service group only.	Do not take the RVG group offline at the primary site.
	■ The RVG group remains online.	
The parent service group fails	■ The child remains online on the primary site.	1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.
	■ An alert notification at the secondary site occurs for the parent only.	2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).
	■ The RVG group remains online.	

Testing fault readiness by running a fire drill

Topics in this chapter include:

- [“About disaster recovery fire drills”](#) on page 985
- [“About the Fire Drill Wizard”](#) on page 986
- [“About post-fire drill scripts”](#) on page 987
- [“Tasks for configuring and running fire drills”](#) on page 988
- [“Prerequisites for a fire drill”](#) on page 988
- [“Fire Drill Wizard actions”](#) on page 990
- [“Preparing the fire drill configuration”](#) on page 992
- [“Running a fire drill”](#) on page 995
- [“Recreating a fire drill configuration that has changed”](#) on page 997
- [“Restoring the fire drill system to a prepared state”](#) on page 998
- [“Deleting the fire drill configuration”](#) on page 1000

About disaster recovery fire drills

A disaster recovery plan should include regular testing of an environment to ensure that a DR solution is effective and ready should disaster strike. This testing is called a fire drill. SFW HA provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. that uses VVR replication.

About the Fire Drill Wizard

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 wizard prepares for the fire drill by completing the following steps:

- Creates a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group, using the same service group name with the prefix *FDnn*. The wizard renames the fire drill service group resources by adding the prefix *FDnn* and changes attribute values as necessary to refer to the FD resources.
- Prepares a copy (mirror) of the production data on the secondary site
You assign one or more disks for the mirrored volumes while running the wizard. Mirror preparation can take some time, so you can exit the wizard once this step is started and let the preparation continue in the background.

Once these steps are complete, the wizard can run the fire drill. Running the fire drill detaches the mirrors from the original volumes to create point-in-time snapshots of the production data. It also brings the application online in the fire drill service group at the secondary site. Bringing the fire drill service group online on the secondary site 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.

Running the fire drill creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes. The wizard assigns the fire drill disk group name by prefixing the original disk group name with *FDnn*.

After you complete the fire drill, it is important to use 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.

Warning: If the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. Therefore, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible after completing the fire drill testing for a service group. See [“Restoring the fire drill system to a prepared state”](#) on page 998.

You must also restore the fire drill configuration before you can delete it. If any errors occur while deleting the fire drill configuration, the wizard will list any incomplete steps so that you can complete them manually.

About post-fire drill scripts

You can specify a script to be run on the secondary site at the end of the fire drill.

For example, if you earlier created and populated a test table at the primary site, you could create a script to verify replication of the data. 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. To run a cmdlet, create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\PowerShell.exe -command "$ScriptName"
```

Where

ScriptName = .ps1 script (fully qualified) / cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\PowerShell.exe -command C:\myTest.ps1
```

Specify the name of the .bat file as the script to run.

Tasks for configuring and running fire drills

The Fire Drill Wizard helps you configure and run a fire drill.

[Table 18-1](#) outlines the high-level objectives and the tasks to complete each objective.

Table 18-1 Tasks for configuring and running fire drills

Objective	Tasks
“Prerequisites for a fire drill” on page 988	Verifying hardware and software prerequisites
“Preparing the fire drill configuration” on page 992	Using the wizard to prepare the initial fire drill configuration
“Running a fire drill” on page 995	Using the wizard to run the fire drill Performing your own tests of the application to confirm that it is operational Note: As soon as you complete the fire drill testing, you should use the wizard to restore the fire drill system. Restoring the fire drill system takes the fire drill service group offline.
“Restoring the fire drill system to a prepared state” on page 998	Using 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.
“Deleting the fire drill configuration” on page 1000	Using the wizard to remove the fire drill configuration

Prerequisites for a fire drill

Ensure that the following prerequisites are met before configuring and running a fire drill:

- The primary and secondary sites must be fully configured with VVR replication and the global cluster option.
- The Veritas FlashSnap option must be installed on all nodes of the clusters at the primary and secondary sites.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.

- All disk groups in the service group are configured for replication. The Fire Drill wizard does not support a configuration in which disk groups are excluded from replication. However, you can exclude individual volumes within a disk group from replication.
- 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.

- For each IP address in the application service group, an IP address must be available to use on the secondary site for the fire drill service group. The wizard can accept input for one IP address and Lanman resource. If the application service group has multiple IP addresses and Lanman resources, the wizard notifies you to edit the fire drill service group resources to supply these values. Information on editing service group resources is covered in the VCS administration guide.
See Veritas Cluster Server Administrator's Guide.
- If you want the fire drill wizard to run a script that you supply, ensure that the script file is available on any secondary site nodes where you plan to run the fire drill.
- 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.

In addition, for testing purposes, you may want to create and populate a new table from the active node at the primary site. After you run the fire drill to bring the fire drill service group online and create the fire drill snapshots, you can check that the table and its data were replicated and are available from the fire drill service group. You can automate this process with a script and when preparing to run the fire drill, specify it as a post-fire drill script.

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.

Fire Drill Wizard actions

While running the Fire Drill Wizard, you select from a menu of fire drill wizard actions on the Fire Drill Mode Selection panel. The actions are listed in the order they are usually performed.

After an action is complete, if you proceed in the wizard, the action menu is displayed again. You can select the next action or exit the wizard and perform the next action later.

The actions consist of the following:

Prepare for Fire Drill	<p>Creates the configuration required to run a fire drill. This step can take some time since the wizard prepares the mirrors for the snapshots.</p> <p>If this option is unavailable, a fire drill configuration already exists on the specified system.</p> <p>See “Preparing the fire drill configuration” on page 992.</p> <p>This option is also displayed while the wizard is recreating a fire drill configuration that has changed.</p> <p>See “Recreating a fire drill configuration that has changed” on page 997.</p>
------------------------	--

Run Fire Drill	<p>Runs the fire drill. The wizard creates the volume snapshots and brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>If a fire drill has been run, you must restore the fire drill configuration to a prepared state before the wizard re-enables this option.</p> <p>See “Running a fire drill” on page 995.</p> <p>Note: If a fire drill configuration exists for the selected service group, the wizard checks for differences in resource names. If differences are found, the wizard can recreate the fire drill configuration before running the fire drill.</p> <p>See “Recreating a fire drill configuration that has changed” on page 997.</p> <p>Note: After completing the fire drill testing, run the wizard again as soon as possible and select the option to restore the configuration. Otherwise the fire drill service group remain online. Be sure to restore one fire drill service group to a prepared state before running a fire drill on another service group.</p>
Restore to Prepared State	<p>Restores the fire drill configuration for another fire drill or to prepare the fire drill configuration for deletion.</p> <p>This option becomes available once a fire drill has been run.</p> <p>The wizard snaps back the snapshot mirrors to reattach to the original volumes and takes the fire drill service group offline.</p> <p>See “Restoring the fire drill system to a prepared state” on page 998.</p>
Delete Fire Drill Configuration	<p>Deletes the fire drill configuration to free up disk space. The wizard deletes the service group on the secondary site and performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill.</p> <p>If a fire drill has been run, this option is disabled until you first restore the fire drill configuration to a prepared state. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p> <p>See “Deleting the fire drill configuration” on page 1000.</p> <p>This option is also displayed while the wizard is recreating a fire drill configuration that has changed.</p> <p>See “Recreating a fire drill configuration that has changed” on page 997</p>

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group and snapshot mirrors of production data at the specified node on the secondary site. 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.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration, you should verify that you meet the prerequisites.

See [“Prerequisites for a fire drill”](#) on page 988.

To prepare for the fire drill

- 1 Open the Solutions Configuration Center (**Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**).
- 2 Start the Fire Drill Wizard (expand **Solutions for Microsoft SQL**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
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.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**. (You can select only one service group at a time for a fire drill.)
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.
The selected system must have access to the replicated data and to disks for the snapshots that will be created for the fire drill.
- 7 If the Recreate Fire Drill Configuration panel is displayed, a fire drill configuration already exists but is not up to date. You can choose whether to run a fire drill or to recreate the configuration first to bring it up to date. If this panel is displayed, see the following procedures:
 - [“Recreating a fire drill configuration that has changed”](#) on page 997.
 - [“Running a fire drill”](#) on page 995Otherwise, continue with the next step.

- 8** In the Fire Drill Mode Selection panel, the available options depend on whether the fire drill service group already exists on this system and whether it is online or offline. Choose one of the following and click **Next**:

If the Prepare for Fire Drill option is available, a fire drill service group does not exist on this system. Click **Prepare for Fire Drill** and continue with the remaining steps in this procedure.

If the Run Fire Drill option is available, a fire drill service group has already been prepared and is up to date. 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 995.

If the Restore to Prepared State option is available, the fire drill service group remains online from a previous fire drill. Click **Restore to Prepared State** and follow the procedure for restoring the fire drill configuration to a prepared state. See [“Restoring the fire drill system to a prepared state”](#) on page 998.

- 9** In the Fire Drill Service Group Settings panel, assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. These must be an address and name not currently in use.

If the service group contains more than one IP and Lanman resource, this panel does not display. After the fire drill service group is created, the wizard notifies you to manually update the IP and Lanman resources in the fire drill service group.

- 10** In the Disk Selection panel, review the information and make the selections as follows and click **Next**:

Volume Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message.

Note: The Disk Selection panel also appears if the wizard is recreating a fire drill service group to which volumes have been added. In that case, only the new volumes are shown for selection.

Disk Group Shows the name of the disk group that contains the original volumes. This field is display only.

Fire Drill DG	Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with <i>FDnn</i> .
Disk	<p>Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume.</p> <p>You can store multiple snapshot volumes on the same disk, if the production volumes reside on disks in the same disk group.</p> <p>If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.</p>
Mount Details	Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This field is display only.

- 11 Wait while the wizard completes the preparation tasks. First the fire drill service group is created on the secondary site (but remains offline). Next 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.
- 12 To run the fire drill now, choose **Next**, or click **Finish** to exit the wizard. If you choose **Finish** the fire drill preparation remains in place. The next time you run the wizard, you choose the primary and secondary systems and service group and then can continue with running the fire drill.

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
 - Splits the fire drill disk group
 - Enables the firedrill resources
 - Brings the fire drill service group online
 - Optionally, executes a specified command to run a script
- For example, if you earlier created and populated a test table at the primary site, you could create a script to verify replication of the data. For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.
- Optionally you can specify a Windows PowerShell cmdlet rather than a script.

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.

If a fire drill has been run previously, you must restore the fire drill configuration to a prepared state before the wizard enables the option to run another fire drill.

Warning: After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, remember to take the fire drill service group offline as soon as possible. You do this by running the wizard and selecting the option to restore the system to the prepared state. 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 998.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, or if you are returning to this procedure after recreating a fire drill service group, go to [step 8](#). Otherwise, if you need to restart the wizard, continue with the next step.

- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft SQL**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
- 5 In the Service Group Selection panel, select the service group and click **Next**.
- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site and click **Next**.
If the fire drill configuration is in a prepared state, the wizard compares the resources of the fire drill service group with the resources of the application service group.
- 7 If the application service group changed since the fire drill configuration was prepared, the wizard displays the Recreate Fire Drill Service Group panel, showing the differences. Choose one of the following:
 - Leave the option checked to recreate the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.
For more information, see [“Recreating a fire drill configuration that has changed”](#) on page 997.
 - Clear the option to recreate the fire drill service group and click **Next**.
- 8 In the Fire Drill Mode Selection panel, click **Run Fire Drill** and click **Next**.
- 9 In the Post Fire Drill Script panel, optionally specify the full path to a script for the wizard to run on the secondary system right after running the fire drill. The script must already exist on the secondary system. Click **Next**.
For more information on post fire drill scripts, see [“About post-fire drill scripts”](#) on page 987.
- 10 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and the Fire drill ran successfully message is displayed.
- 11 Click **Finish**.
- 12 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.

- 13 Run the wizard again to restore the fire drill configuration to the prepared state.
See [“Restoring the fire drill system to a prepared state”](#) on page 998.

Recreating a fire drill configuration that has changed

When you run the Fire Drill wizard, a fire drill service group may already exist for the selected application service group. However, the application service group may have changed since the fire drill service group was created. Therefore, the wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Recreate Fire Drill Service Group panel.

You have the following choices from the Recreate Fire Drill Service Group panel:

- Leave the option checked to recreate the fire drill service group. Proceed with using the wizard to recreate the configuration to match the application service group. The wizard deletes the fire drill configuration but does not delete the mirrors for volumes that still exist. It then recreates the fire drill configuration, preparing new mirrors only for new volumes. If volumes have been removed, the wizard displays an additional option to snap abort the obsolete snapshot volumes to free up disk space.
- Clear the option to recreate the fire drill service group. You can then proceed with using the wizard to do either of the following:
 - Run the fire drill, ignoring the differences.
 - Delete the entire fire drill configuration. Then start over with preparing the fire drill configuration.

Note: The wizard does not check for changes in volume attributes, such as the MountPath attribute. For example, if you have a MountV resource with an attribute that points to drive Y and you change that attribute to point to drive X, the wizard does not identify this change and does not give the option to recreate the fire drill service group.

You can choose whether to manually edit the fire drill service group for such changes and then run the fire drill, ignore the differences, or delete the configuration and start over.

The following procedure describes the choice of recreating the fire drill configuration.

To recreate the fire drill configuration if the service group has changed

- 1 In the Recreate Fire Drill Service Group panel, leave the option checked to recreate the configuration before running the fire drill. 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. 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. When all tasks are complete, click **Next**.

Warning: If you close the wizard after deleting the fire drill configuration without continuing on to the fire drill preparation step, the information on the existing snapshot volumes is lost.

- 4 In the Fire Drill Mode Selection panel, Prepare for Fire Drill is selected. Click **Next**.
- 5 If volumes have been added, the Disk Selection panel is displayed for the new volumes. Specify the information for the added volumes. Click **Next**. 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.
- 6 Wait while the wizard recreates the fire drill service group and 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 mirror preparation is complete, click **Next**. Continue with running the fire drill.
See “[Running a fire drill](#)” on page 995.

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 after you select the Prepare for Fire Drill option.

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.

When restoring the fire drill system to a prepared state, the wizard completes the following tasks:

- Takes the fire drill service group offline
- Disables the fire drill service group resources
- Imports the fire drill disk group
- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to [step 8](#). Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand **Solutions for Microsoft SQL**, 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 Mode Selection panel, click **Restore to Prepared State** and click **Next**. Click **Yes** on the confirmation message.
If you have run a fire drill but not yet restored the configuration, this is the only option available. If the option is unavailable, the configuration has already been restored or the fire drill has not yet been run.
- 8 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next** if you want to delete the fire drill configuration or 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 and 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.

If you have run a fire drill and want to delete the configuration, you must first restore the fire drill configuration to a prepared state before the wizard enables the option to delete the fire drill configuration.

See “[Restoring the fire drill system to a prepared state](#)” on page 998.

To delete a fire drill configuration

- 1 If you have just used the wizard to restore the 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 Microsoft SQL**, expand **Fire Drill**, expand **Configure or run a fire drill**, and click **Fire Drill Wizard**).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 If the wizard detects that the fire drill service group is different from the application service group, it displays the Recreate Fire Drill Service Group panel. Clear the option to recreate the fire drill service group and click **Next**.
- 8 In the Fire Drill Mode Selection panel, click **Delete Fire Drill Configuration** and click **Next**, and click **Yes** to confirm the deletion.
- 9 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.

Appendices

This section includes the following appendices:

- [Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation](#)
- [Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation](#)
- [Configuring the DR components \(VVR and GCO\) without using the DR wizard](#)
- [Configuring an MSDTC service group for disaster recovery](#)

Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation

This appendix includes the following topics:

- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring the cluster](#)
- [Configuring cluster disk groups and volumes](#)
- [Installing and configuring SQL Server 2000 on the first node](#)
- [Preparing to install SQL Server on the second node](#)
- [Installing SQL Server 2000 on the second node](#)
- [Setting the internal name of the clustered instance](#)
- [Configuring the VCS SQL Server service group](#)
- [Creating a SQL Server user-defined database](#)
- [Verifying the cluster configuration](#)

- [Creating a parallel environment on the secondary site](#)
- [Installing DR components on primary and secondary sites](#)

Note: This appendix covers an earlier method of deploying disaster recovery. A newer method uses the Solutions Configuration Center and the Disaster Recovery (DR) wizard to clone storage configuration and service groups and set up replication. See [“Deploying disaster recovery: New SQL Server 2000 installation”](#) on page 779.

After setting up a SFW or SFW HA environment for Microsoft SQL Server on a primary site, you must create a secondary or “failover” site for disaster recovery. This chapter provides information on how to install and configure the high availability and SQL components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171.

To complete the process of creating a DR solution, you must proceed to [“Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 1149 after performing the tasks outlined in this chapter.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation

Objective	Tasks
“Reviewing the prerequisites” on page 1007	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 1011	<ul style="list-style-type: none"> ■ Understanding active-passive configuration and site failover in a DR environment ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 1014	<ul style="list-style-type: none"> ■ Setting up the storage hardware for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation (Continued)

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 1017	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ■ Selecting the Global Cluster Option for VCS to enable wide-area failover ■ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ■ Configuring VxSAS
“Configuring the cluster” on page 1026	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 1043	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2000 on the first node” on page 1051	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2000 ■ Configuring SQL services
“Preparing to install SQL Server on the second node” on page 1054	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters ■ Removing shared SQL files from the cluster disk group

Table A-1 Tasks for manually implementing disaster recovery for a new SQL Server 2000 installation (Continued)

Objective	Tasks
“Installing SQL Server 2000 on the second node” on page 1058	Installing SQL Server 2000
“Setting the internal name of the clustered instance” on page 1062	Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 1064	Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating a SQL Server user-defined database” on page 1069	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server ■ Adding resources for a user-defined database in VCS
“Verifying the cluster configuration” on page 1071	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“Creating a parallel environment on the secondary site” on page 1072	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Configuring the cluster ■ Configuring disk groups and volumes for SQL
“Installing DR components on primary and secondary sites” on page 1073	<ul style="list-style-type: none"> ■ Completing the tasks outlined in “Configuring the DR components (VVR and GCO) without using the DR wizard” on page 1149

Reviewing the prerequisites

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware at:

<http://www.symantec.com/business/support/index.jsp>

Supported Software

- Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL
 - For a disaster recovery installation include the Global Clustering Option and depending on your replication solution, Veritas Volume Replicator or a hardware replication agent

For a Microsoft SQL Server 2000 environment, any of the following SQL Servers and their operating systems:

- | | |
|---|---|
| <p>Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) ■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required for all editions) |
| <p>Microsoft SQL Server 2000 (64-bit) Enterprise Edition</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both) |
| <p>Microsoft SQL Server 2000 (64-bit) Standard Edition or Enterprise Edition (SP4 required)</p> | <ul style="list-style-type: none"> ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required for all editions) ■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required for all editions) |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 1010.
- 1 GB of RAM per server for SFW HA.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL Virtual Server

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

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C : \WINDOWS of one node, installations on all other nodes must be on C : \WINDOWS. Make sure that the

same drive letter is available on all nodes and that the system drive has adequate space for the installation.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Disk space requirements

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

[Table A-2](#) on page 1010 estimates disk space requirements for SFW HA.

Table A-2 Disk space requirements

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

Reviewing the configuration

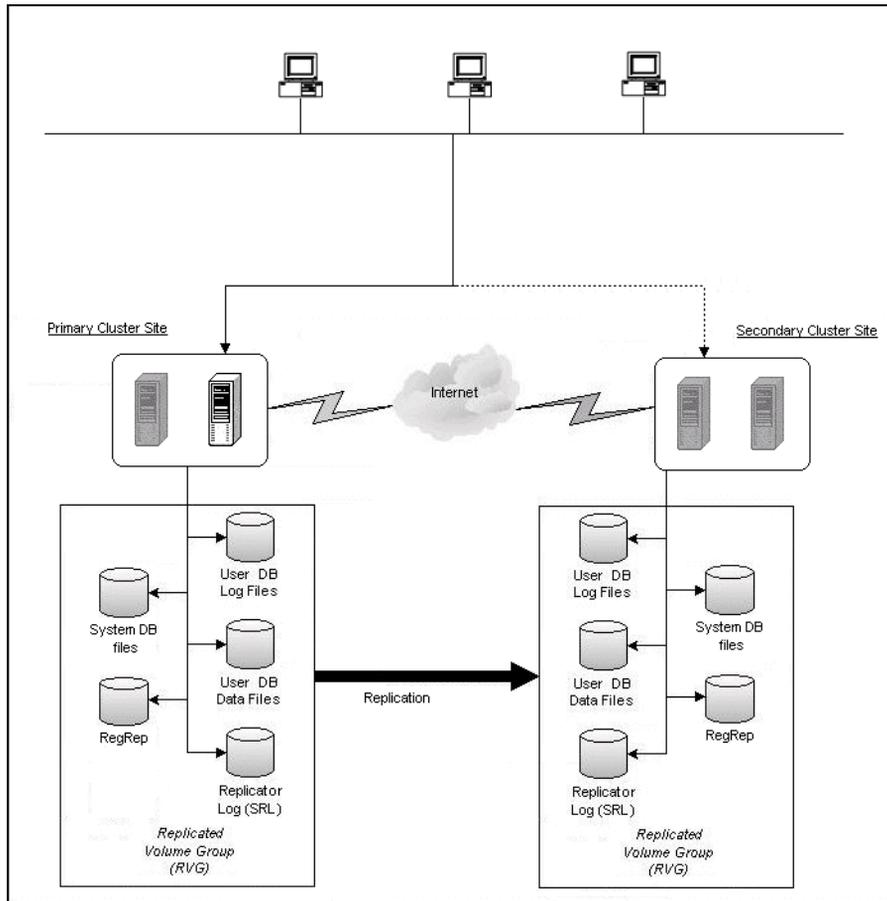
During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under [“Reviewing the prerequisites”](#) on page 1007.

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2000 application data is stored on the volumes that are under the control of the RVG.

Figure A-1 Typical VVR configuration



If the Microsoft SQL Server 2000 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

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:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2000 service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table A-3](#) on page 1017 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table A-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see [“Installing Symantec Trusted certificate for unsigned drivers”](#) on page 1018.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

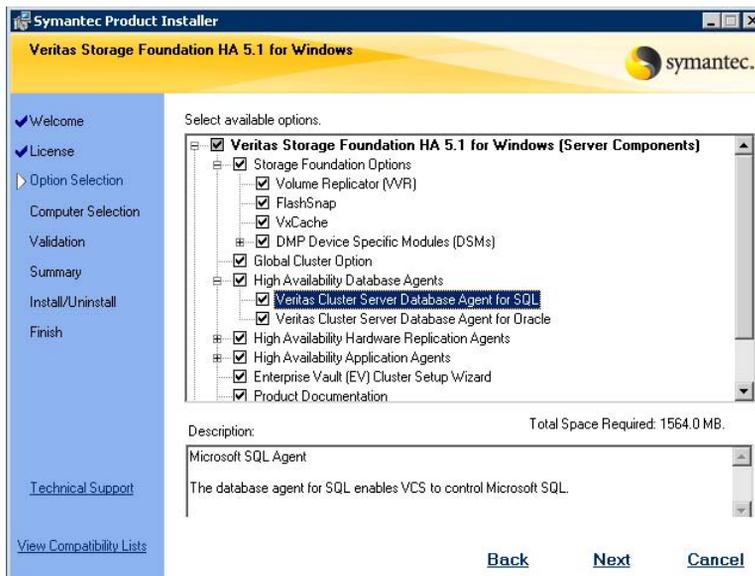
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

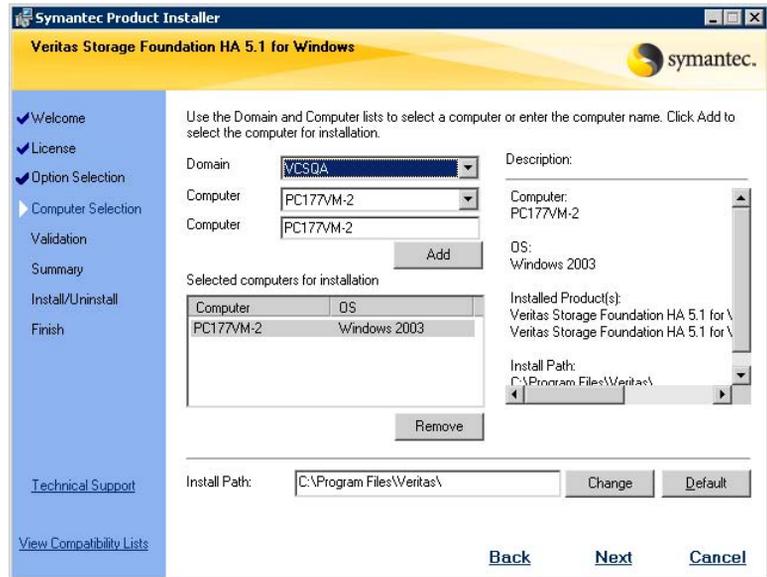
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



- | | |
|---|---|
| Veritas Cluster Server Database Agent for SQL | Required to configure high availability for SQL Server. |
| Client | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration. |
| Global Cluster Option | Required for a disaster recovery configuration only. |
| Veritas Volume Replicator | For a disaster recovery configuration only: if you plan to use VVR for replication, select the option to install VVR. |
| High Availability Hardware Replication Agents | If you plan to use hardware replication, select the appropriate hardware replication agent. |

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Configuring VxSAS

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxscfg.exe` from the command prompt of the required machine.
Read the information provided on the Welcome page and click **Next**.
- 2 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 the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts.

Click **Next**.

- 3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

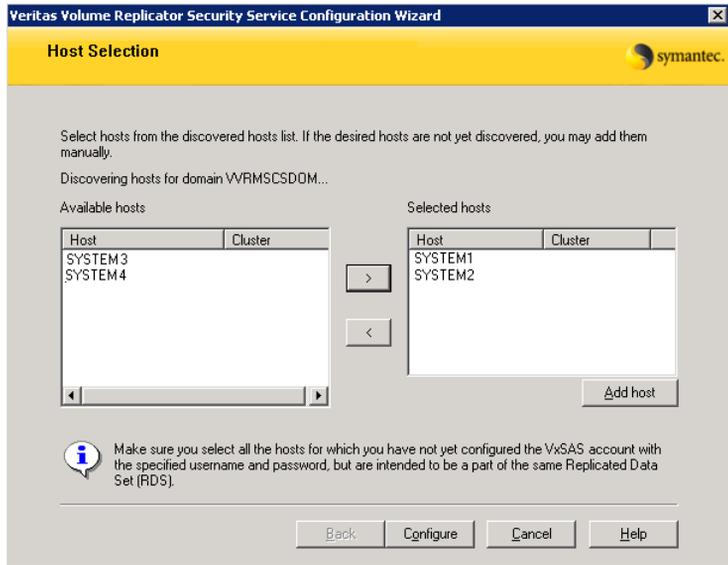
Selecting domains The Available domains pane lists all the domains that are present in the Windows network neighborhood.

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.

Click **Back** to change any information you had provided earlier.

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

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

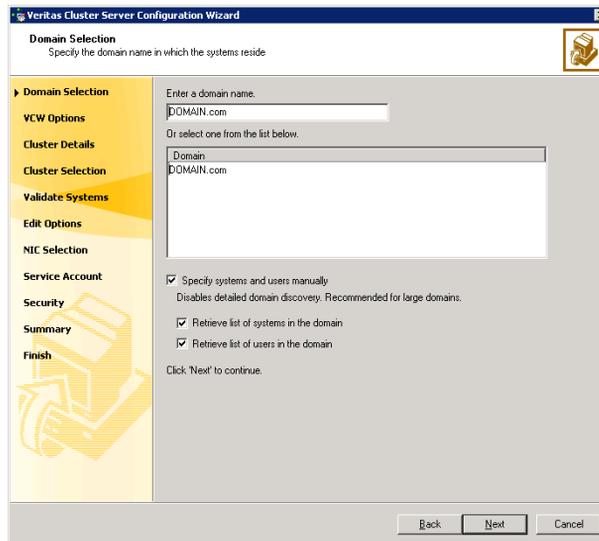
Complete the following tasks before creating a cluster:

- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
 - If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group. If you plan to use an existing user account for the VCS Helper service, you must know the password for the user account.

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

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

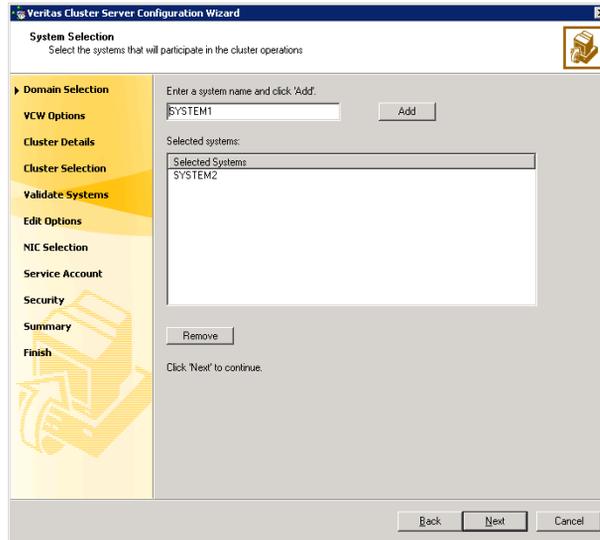
Proceed to [step 8](#) on page 1029.

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

- Check the **Specify systems and users manually** check box.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.

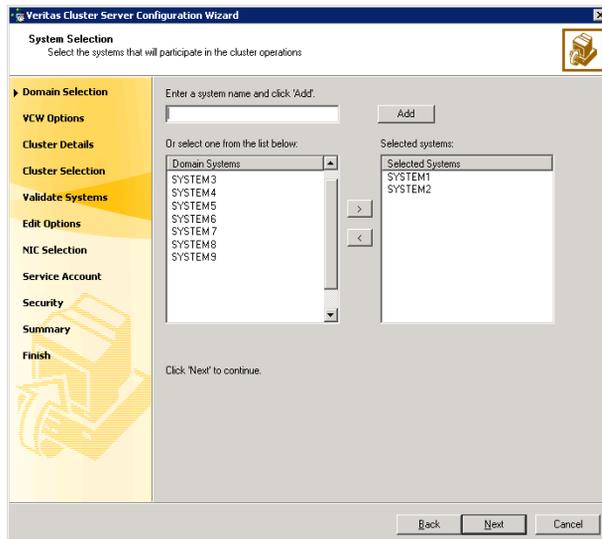
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 1028. Otherwise proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 1029.

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



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

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

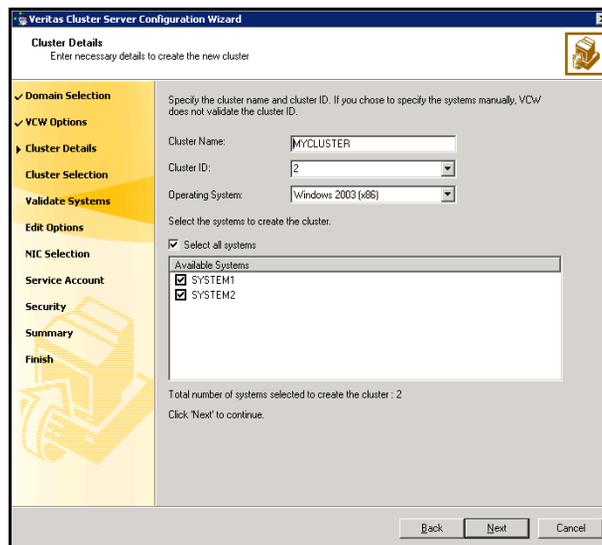
A system can be rejected for any of the following reasons:

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

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.
- 9 On the Cluster Details panel, specify the details for the cluster and then click **Next**.



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 1027 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

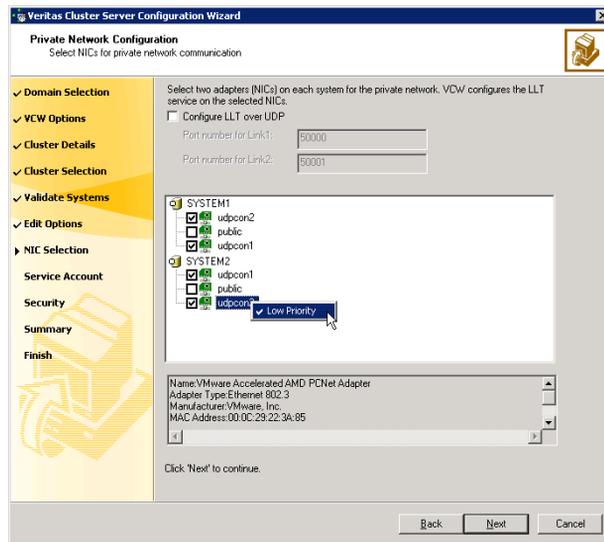
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 1029, proceed to the next step. Otherwise, proceed to [step 12](#) on page 1033.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

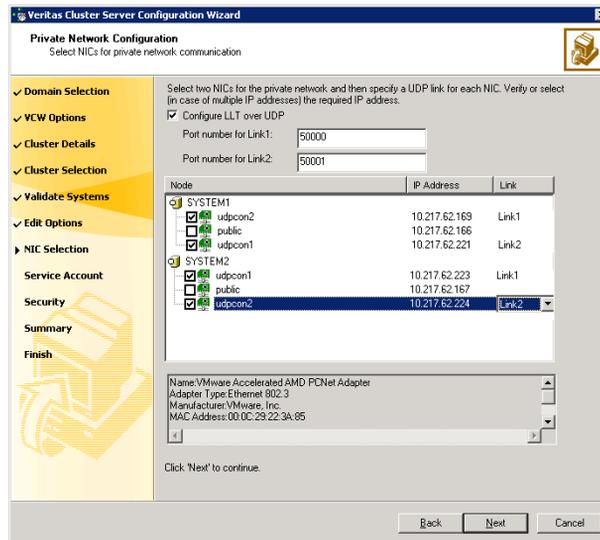
- To configure the VCS private network over Ethernet



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

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

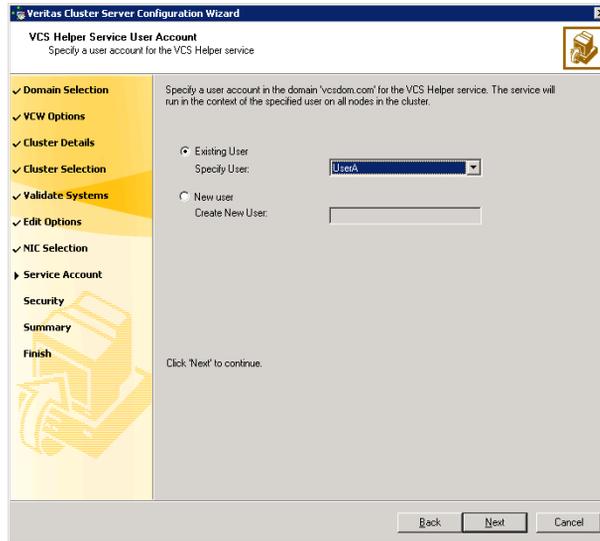
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



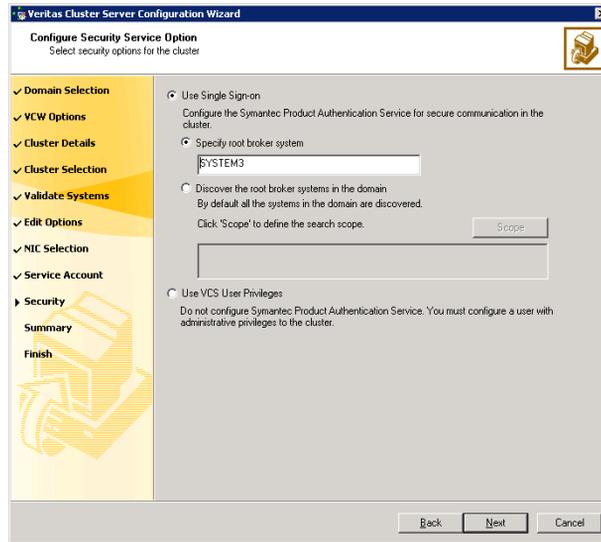
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 1027, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
 Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.

[Table A-4](#) on page 1035 contains some more examples of search criteria.

Table A-4 Search criteria examples

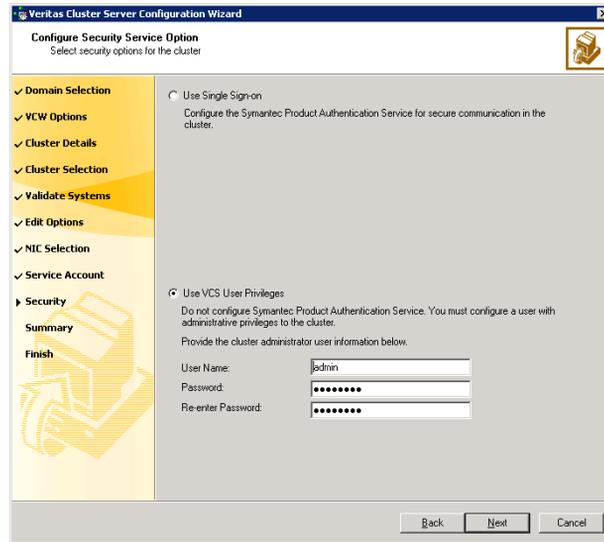
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**.

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

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

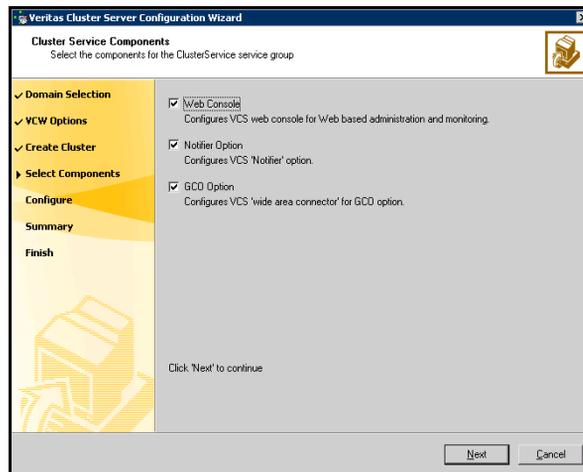
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Web Console, notification, and for global clusters.

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

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

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

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



- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.

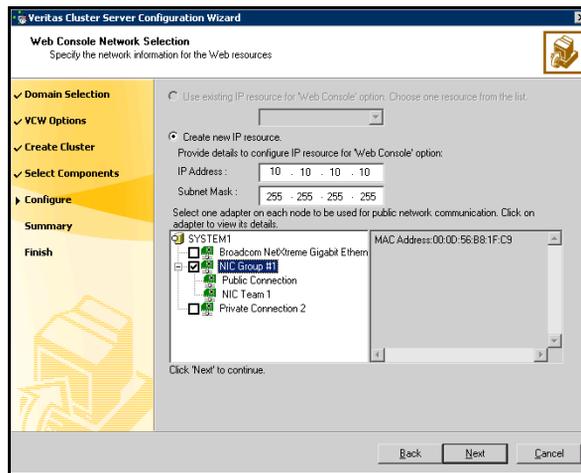
The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option

Configuring the Web Console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



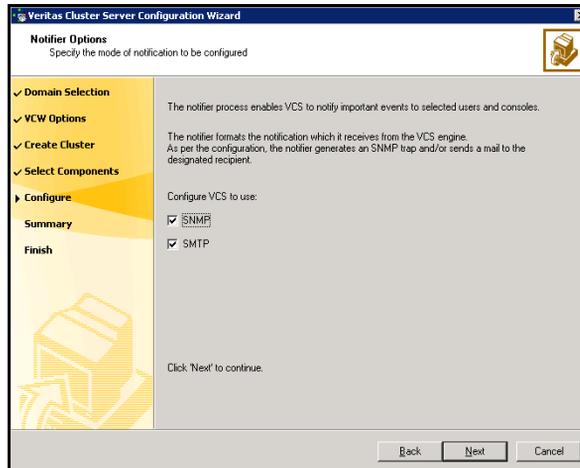
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to “[Configuring notification](#)” on page 1039.
If you chose to configure global cluster components, proceed to “[Configuring the wide-area connector process for global clusters](#)” on page 1042.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

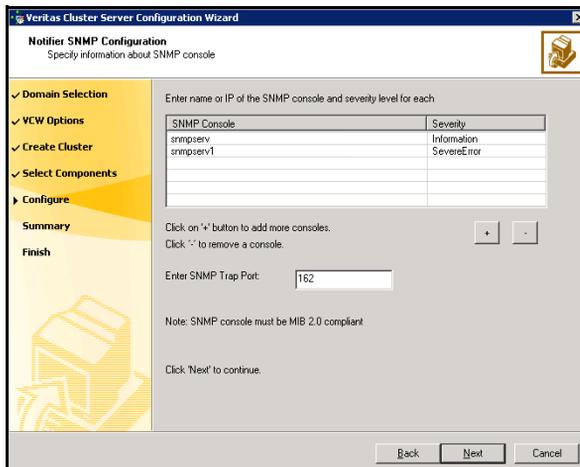
To configure notification

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

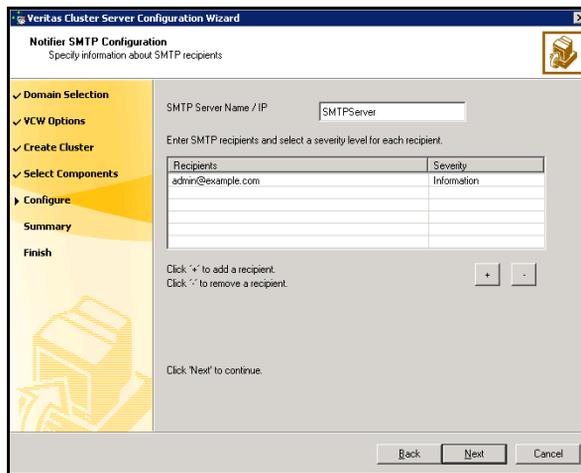


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

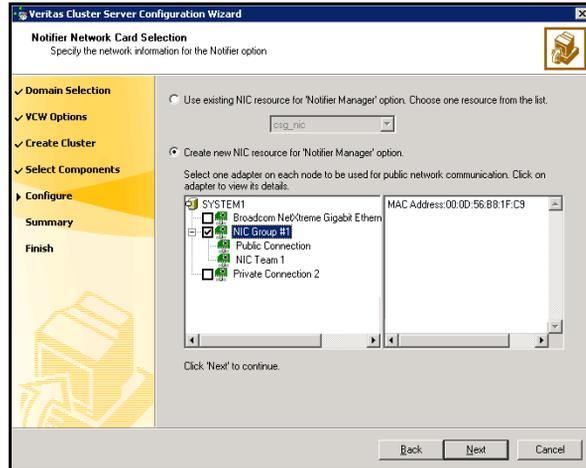


- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click + to add a field; click - to remove a field.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



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

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



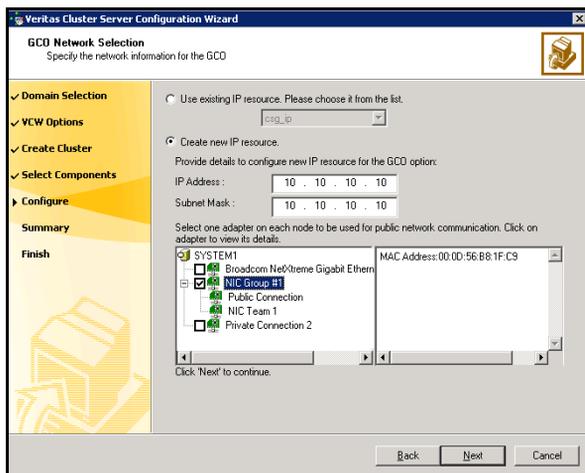
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.
 - 6 If you chose to configure global cluster components, proceed to [“Configuring the wide-area connector process for global clusters”](#) on page 1042. Otherwise, click **Finish** to exit the wizard.

Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
 - 3 Click **Finish** to exit the wizard.

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

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

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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 disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server use database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REPLOG: Contains the VVR Storage Replicator Log.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Use the following procedures to create the appropriate disk groups and volumes. Note that instructions for INST1_REPLOG are provided in “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 1151; you can wait until that chapter to configure VVR and its related volumes.

The first time you open the Veritas Enterprise Administrator (VEA), it will automatically connect to the localhost. It will also display the names of other hosts that are connected to VEA as nodes in the left side tree display.

Creating a cluster disk group

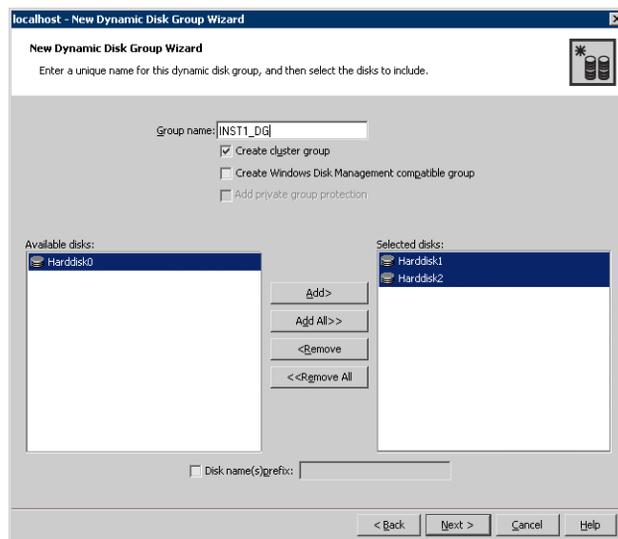
Create a cluster disk group on the first node of the cluster.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_DB1_VOL:** contains the user database files
- **INST1_DB1_LOG:** contains the user database log files
- **INST1_REPLOG:** Contains the VVR Storage Replicator Log.

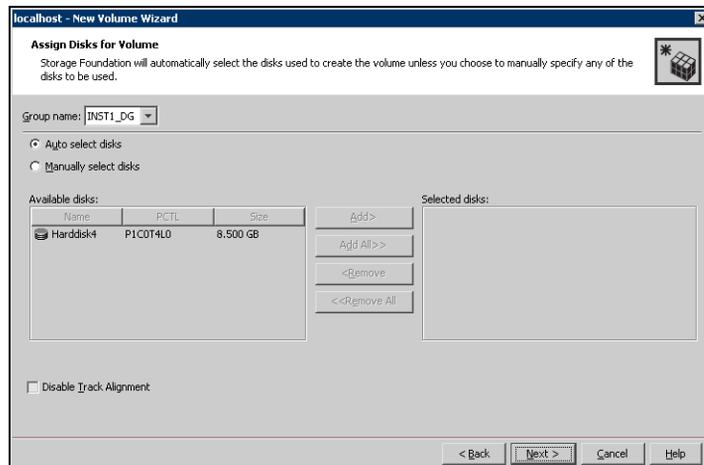
Caution: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

Refer to the steps below for the **INST1_DATA_FILES** and **INST1_REGREP_VOL** volumes. You can create the **INST1_DB1_VOL** and **INST1_DB1_LOG** volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 1069. You can create the **INST1_REPLOG** volume at this time or during the process of “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 1151.

Note: To ensure that the drive letters you assign to the new volumes will always be available on all nodes, assign drive letters starting in the middle of the alphabet. This way when drive letters are assigned as additional internal devices are added to a node there will not be a conflict with the volume drive letters.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
 You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.



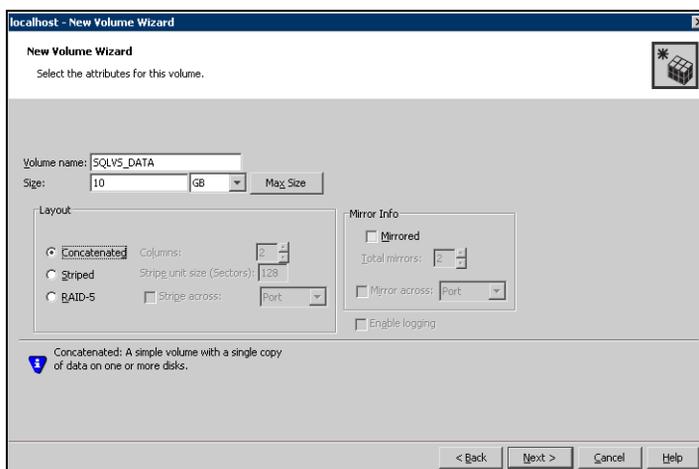
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.

- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

- Click **Next**.

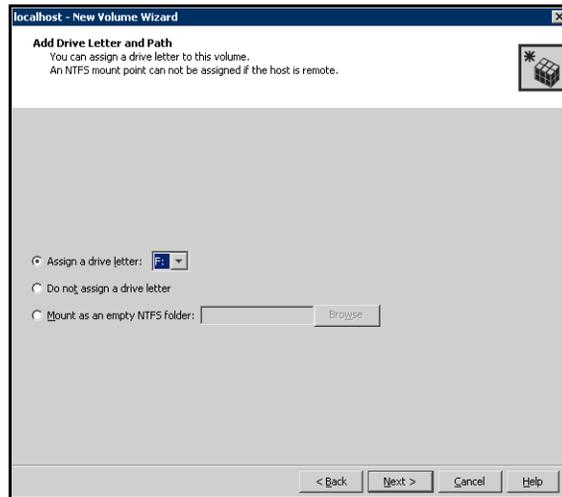
7 Specify the parameters of the volume.



- Enter a name for the volume. 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.
- 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.
- Verify that **Enable Logging** 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

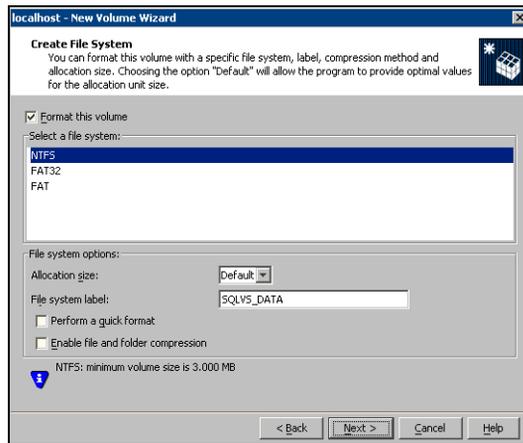
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- 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.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - 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**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Installing and configuring SQL Server 2000 on the first node

In preparation for installing Microsoft SQL Server 2000, ensure the cluster disk group is imported to the first node and the volumes are mounted. Complete the following procedures to install and configure Microsoft SQL Server 2000:

- [Installing Microsoft SQL Server](#)
- [Setting SQL Server 2000 services to manual start](#)

Installing Microsoft SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the first node and the volumes are mounted (are assigned drive letters) See “[Importing the cluster disk group](#)” on page 1056 and “[Adding drive letters to mount the volumes](#)” on page 1056.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.

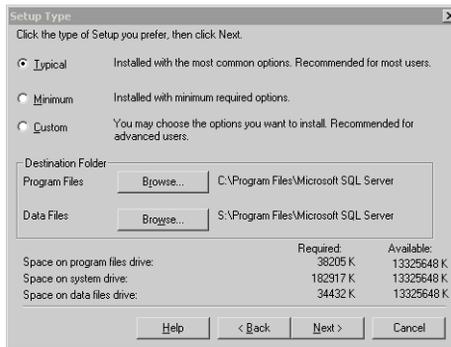
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

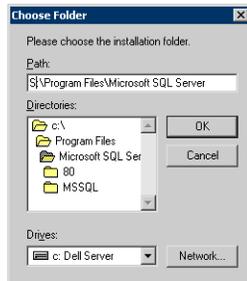
To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.

- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.

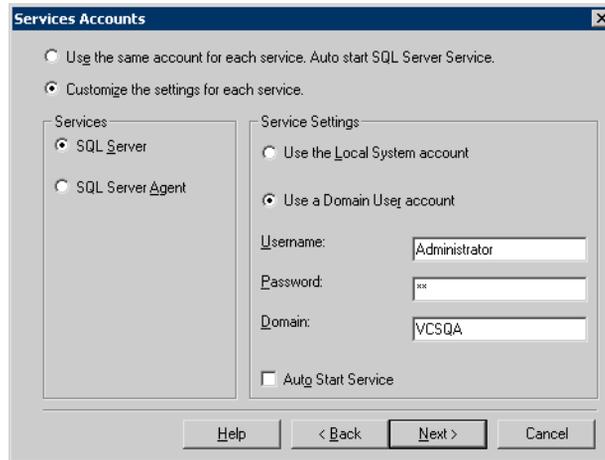


- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

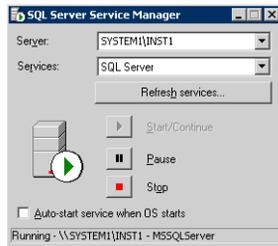
Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Setting SQL Server 2000 services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - Select a service from the **Services** list.
 - Clear the **Auto-start service when OS starts** check box.
- 2 Repeat these steps for all other SQL Server services that are running on the server.

Preparing to install SQL Server on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2000 Service”](#) on page 1055
- [“Deporting the cluster disk group”](#) on page 1055
- [“Importing the cluster disk group”](#) on page 1056
- [“Adding drive letters to mount the volumes”](#) on page 1056
- [“Renaming shared SQL Server 2000 files”](#) on page 1058

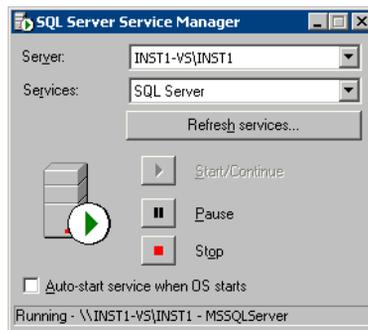
Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2000 Service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Deporting the cluster disk group

In order to install SQL Server 2000 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).

- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

- 1 In the Actions menu, select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - Click **OK** in the Import Dynamic Disk Group dialog box.

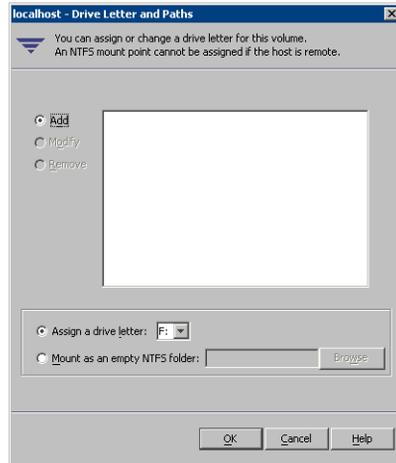
Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.

- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 1058
- [“Removing shared SQL Server files”](#) on page 1061

Installing SQL Server

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported to the additional node and the volumes are mounted (are assigned drive letters). See [“Importing the cluster disk group”](#) on page 1056 and [“Adding drive letters to mount the volumes”](#) on page 1056.

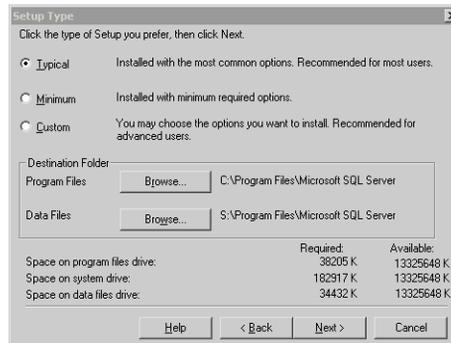
Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

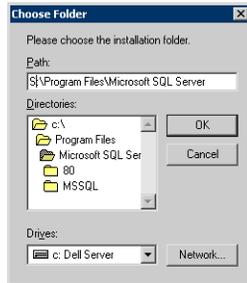
Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.

To install Microsoft SQL Server 2000

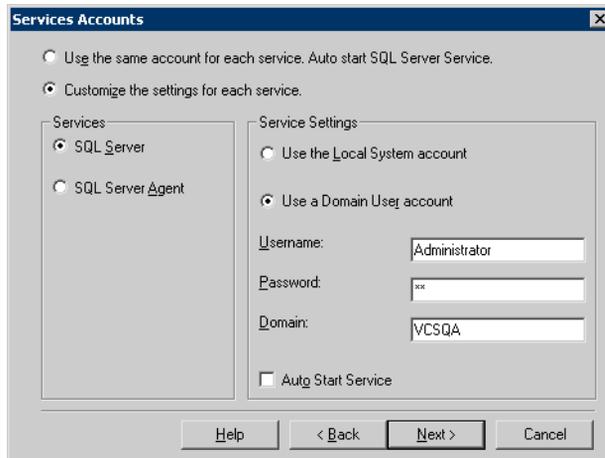
- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
 Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.
- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.



- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
 - For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.
- 10 In the Service Accounts panel, make the following selections and click **Next**:



- Choose the **Customize the settings for each service** option.
 - In the Services box, select the **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and then specify the user name, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation.

- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Repeat the procedures described in “[Preparing to install SQL Server on the second node](#)” on page 1054 and “[Installing SQL Server 2000 on the second node](#)” on page 1058 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

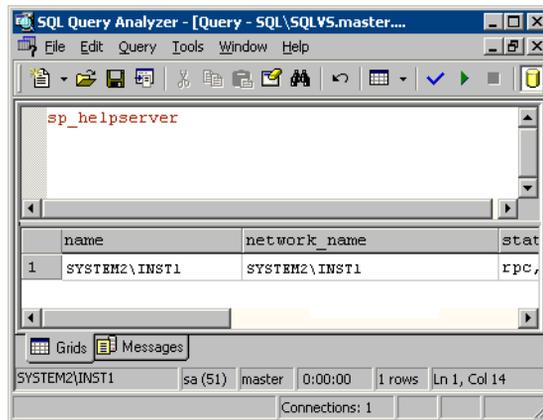
To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped** checkbox.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
`“sp_dropserver ‘System_Name\Instance_Name.’”`
 where *System_Name\Instance_Name* is the name noted in step 3.
 For example, for named instance:
`“sp_dropserver ‘SYSTEM2\INST1.’”`
 For example, for a default instance:
`“sp_dropserver ‘SYSTEM1.’”`
 - Press F5.
- 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter
`"sp_addserver 'Virtual_Server_Name\Instance_Name', local"`
For example `'INST1-VS\INST1'`, `local` for a named instance, or `'INST1-VS'`, `local` for a default instance.
 - Press F5.

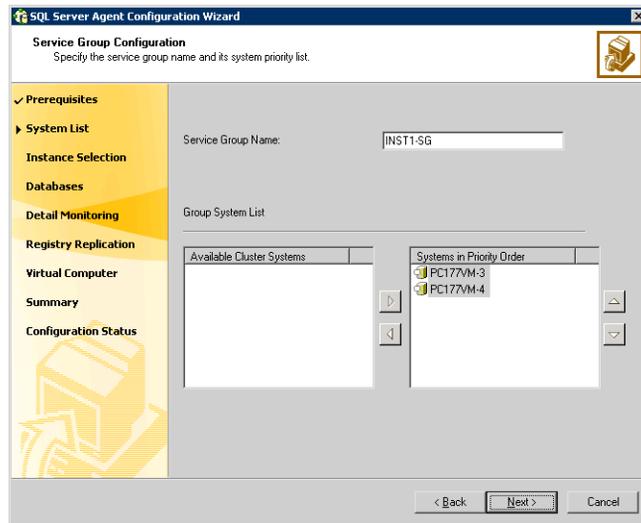
Configuring the VCS SQL Server service group

The Enterprise Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

- Verify that SFW HA, along with the VCS enterprise agent for SQL Server 2000, is installed on all cluster nodes. See ["Configuring the cluster"](#) on page 1026.
- Verify you have configured a VCS cluster using VCS Configuration Wizard (VCW). See ["Configuring the cluster"](#) on page 1026.
- Verify you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify the drive containing the SQL Server 2000 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify the SQL Server 2000 instance is installed identically on all nodes that will participate in the service group.
- Stop the SQL 2000 Server service for the SQL instance. See ["Stopping the SQL Server 2000 Service"](#) on page 1055.
- Assign a unique virtual server name to SQL Server 2000. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary site.
- Assign a unique virtual IP address to the SQL Server 2000 instance.

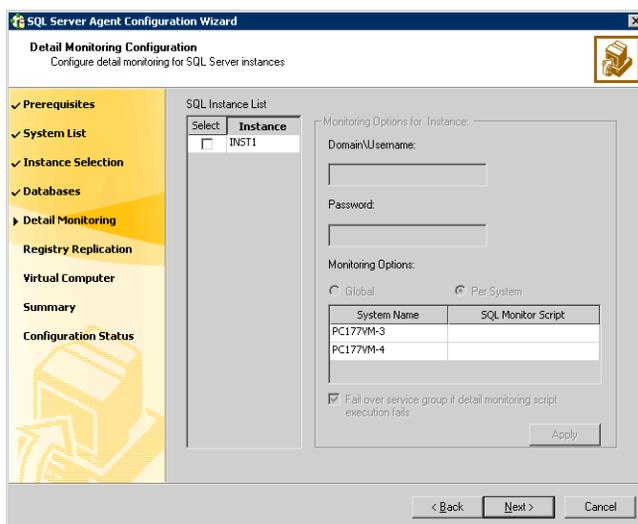
To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

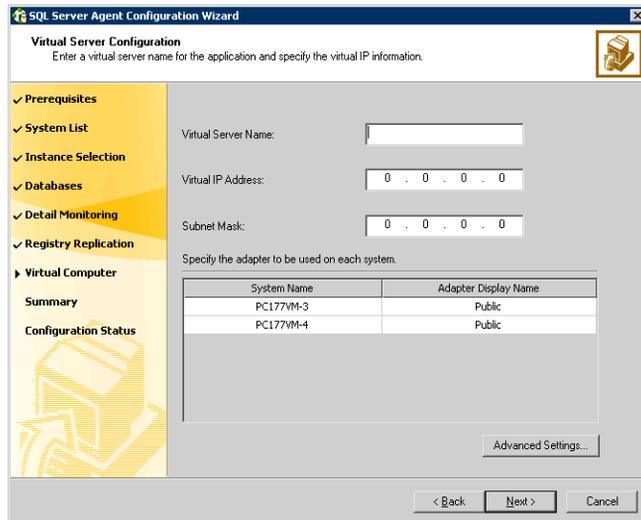
- Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.
Click **Next**.
- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 1066 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to

fail over the service group if the detail monitoring script execution fails.

- Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 1066 for all the instances for which detail monitoring will be configured.
- 10 Click **Next**.
- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

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

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 1171.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 1046 for information on how to use the VEA console to create a volume.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and select **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the **Location** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click the browse button (...) in the **Location** column and browse to the location of the volume you created for the transaction log, and click **OK**.

Adding VMDg and MountV resources

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for disaster recovery](#)” on page 1171.

Verifying the cluster configuration

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

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

To shut down an active cluster node

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

Creating a parallel environment on the secondary site

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

- [“Reviewing the prerequisites”](#) on page 1007
- [“Reviewing the configuration”](#) on page 1011
- [“Configuring the storage hardware and network”](#) on page 1014
- [“Configuring the cluster”](#) on page 1026
- [“Configuring cluster disk groups and volumes”](#) on page 1043
During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 1051
Select the same options at the secondary site as you did at the primary site.
- [“Preparing to install SQL Server on the second node”](#) on page 1054
The instance name must be the same on the primary site and secondary site.
- [“Installing SQL Server 2000 on the second node”](#) on page 1058
- [“Configuring the VCS SQL Server service group”](#) on page 1064
The service group name, virtual computer name, and IP address must be the same on both the primary site and secondary site.

Note: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

Installing DR components on primary and secondary sites

To complete the process of creating a DR solution, proceed to [Appendix C, “Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 1149 after performing the tasks outlined in this chapter.

1074 | Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation
Installing DR components on primary and secondary sites

Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation

This appendix includes the following topics:

- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring the storage hardware and network](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring the cluster](#)
- [Configuring cluster disk groups and volumes](#)
- [Installing and configuring SQL Server 2005 on the first node](#)
- [Preparing to install SQL Server 2005 on the second node](#)
- [Installing SQL Server 2005 on the second node](#)
- [Setting the internal name of the clustered instance](#)
- [Configuring the VCS SQL Server service group](#)
- [Creating a SQL Server user-defined database](#)
- [Verifying the cluster configuration](#)

- [Creating a parallel environment on the secondary site](#)
- [Installing DR components on the primary and secondary sites](#)

Note: This appendix covers an earlier method of deploying disaster recovery. You can now use the Solutions Configuration Center and the Disaster Recovery (DR) wizard to clone storage configuration and service groups and assist you in setting up replication. See [“Deploying disaster recovery: New SQL Server 2005 installation”](#) on page 881.

After setting up a SFW or SFW HA environment for Microsoft SQL Server on a primary site, you must create a secondary or “failover” site for disaster recovery. This chapter provides information on installing and configuring the high availability and SQL components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

To configure MSDTC service groups, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171. To complete the process of creating a DR solution, you must proceed to [“Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 1149 after performing the tasks outlined in this chapter.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation

Objective	Tasks
“Reviewing the prerequisites” on page 1079	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 1084	<ul style="list-style-type: none"> ■ Understanding active-passive configuration and site failover in a DR environment ■ Reviewing the sample configuration
“Configuring the storage hardware and network” on page 1087	<ul style="list-style-type: none"> ■ Setting up the storage hardware and network for a cluster environment ■ Verifying the DNS entries for the systems on which SQL will be installed

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation (Continued)

Objective	Tasks
“Installing Veritas Storage Foundation HA for Windows” on page 1090	<ul style="list-style-type: none"> ■ Verifying the driver signing option for the system ■ Installing Veritas Storage Foundation for Windows HA (automatic installation) ■ Selecting the option to install VVR; this will also automatically install the Veritas Cluster Server Agent for VVR ■ Selecting the Global Cluster Option for VCS to enable wide-area failover ■ Selecting the option to install Veritas Cluster Server Agent for Microsoft SQL Server ■ Configuring VxSAS
“Configuring the cluster” on page 1099	<ul style="list-style-type: none"> ■ Verifying static IP addresses and name resolution configured for each node ■ Configuring cluster components using the VCS Cluster Configuration Wizard (VCW) ■ Setting up secure communication for the cluster
“Configuring cluster disk groups and volumes” on page 1116	<ul style="list-style-type: none"> ■ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator
“Installing and configuring SQL Server 2005 on the first node” on page 1124	<ul style="list-style-type: none"> ■ Installing and configuring SQL Server 2005 ■ Configuring SQL services
“Preparing to install SQL Server 2005 on the second node” on page 1128	<ul style="list-style-type: none"> ■ Stopping the SQL Service ■ Deporting the cluster disk group from the first node ■ Importing the cluster disk group on an additional node ■ Adding drive letters ■ Removing shared SQL files from the cluster disk group

Table B-1 Tasks for manually implementing disaster recovery for a new SQL Server 2005 installation (Continued)

Objective	Tasks
“Installing SQL Server 2005 on the second node” on page 1131	Installing SQL Server 2005
“Setting the internal name of the clustered instance” on page 1135	Setting the internal name of the clustered instance
“Configuring the VCS SQL Server service group” on page 1138	Creating a SQL Server service group using the VCS SQL Configuration Wizard
“Creating a SQL Server user-defined database” on page 1143	<ul style="list-style-type: none"> ■ Creating volumes for a user-defined database and transaction log ■ Creating a new user-defined database in SQL Server ■ Adding resources for a user-defined database in VCS
“Verifying the cluster configuration” on page 1145	<ul style="list-style-type: none"> ■ Simulating failover ■ Switching online nodes
“Creating a parallel environment on the secondary site” on page 1146	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the configuration ■ Configuring the network and storage ■ Installing SFW HA ■ Configuring the cluster ■ Configuring disk groups and volumes for SQL
“Installing DR components on the primary and secondary sites” on page 1147	Completing the tasks outlined in “Configuring the DR components (VVR and GCO) without using the DR wizard” on page 1149

Reviewing the prerequisites

Verify the requirements for your configuration are met before starting the Veritas Storage Foundation HA for Windows installation. This replication recovery solution requires installation and configuration at a primary site and a secondary site.

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

Before you install SFW HA, verify that your configuration meets the following criteria and that you have reviewed the SFW 5.1 Hardware Compatibility List to confirm supported hardware:

<http://www.symantec.com/business/support/index.jsp>

For a Disaster Recovery configuration select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.

Supported software

Microsoft SQL Server

For Microsoft SQL Server, you need Veritas Storage Foundation HA 5.1 for Windows (SFW HA) with the Veritas Cluster Server Database Agent for Microsoft SQL, and any of the following SQL Server environments with the corresponding operating system.

For a Microsoft SQL Server 2005 environment, any of the following SQL Servers and their operating systems:

- | | |
|--|--|
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition or Datacenter Edition (SP2 required)■ Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP2 required)■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2003 (SQL Server 2005 SP1 or higher required) | <ul style="list-style-type: none">■ Windows Server 2003 for Itanium-based Systems Enterprise Edition or Datacenter Edition (SP2 required for both)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required) |
| Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition on Windows Server 2008 (SQL Server 2005 SP2 required) | <ul style="list-style-type: none">■ Windows Server 2008 for 64-bit Itanium (IA64)■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition |

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one "default" instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

System requirements

Systems must meet the following requirements:

- Memory: minimum 1 GB of RAM per server for SFW HA.
- Memory: minimum 1 GB of RAM per server for SQL Server 2005; refer to the Microsoft documentation for more information.
- 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.

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). See the *Veritas Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network. Symantec recommends three NICs. See "[Best practices](#)" on page 1083.
- All servers must run the same operating system, service pack level, and system architecture.

Network requirements

This section lists the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW

HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.

- Static IP addresses for the following purposes:
 - One static IP address available per site for each SQL 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 the following options: Notification, Cluster Management Console (web console), or Global Cluster Option. The same IP address may be used for all options.
 - For VVR only, a minimum of one static IP address per site for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the `DNSRefreshInterval` attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.

See the *Veritas Cluster Server Bundled Agents Reference Guide*.

Permission requirements

This section lists the following permission requirements:

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

Additional requirements

Please review the following additional requirements:

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system.

Best practices

Symantec recommends that you perform the following tasks:

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table B-2](#) on page 1084 estimates disk space requirements for SFW HA.

Table B-2 Disk space requirements

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

Reviewing the configuration

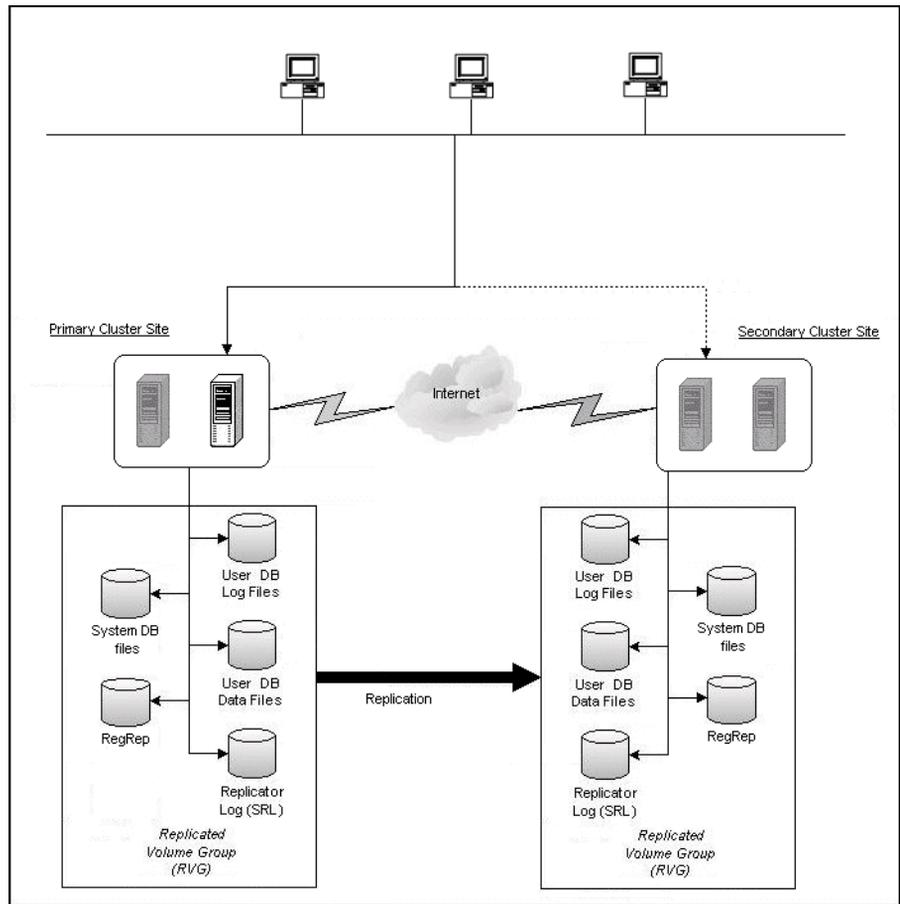
During the configuration process you will create virtual IP addresses.

The virtual IP address for the SQL virtual server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.

For additional IP addresses required, see the “Network requirements” section under “[Reviewing the prerequisites](#)” on page 1079.

The following figure illustrates a typical clustered VVR configuration. In this case the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by VVR for setting up the RVG. The Microsoft SQL Server 2005 application data is stored on the volumes that are under the control of the RVG.

Figure B-1 Typical VVR configuration



If the Microsoft SQL Server 2005 server on SYSTEM1 fails, SQL Server comes online on node SYSTEM2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

Sample configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site. 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:

Primary Site

SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server 2005 service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_REPLOG	replicator log volume required by VVR
INST1	SQL Instance Name

Secondary Site

SYSTEM3 & SYSTEM4	first and second nodes of the secondary site
	All the other parameters are the same as on the primary site.

DR Components

INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Configuring the storage hardware and network

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

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

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on each system.
 - To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Symantec recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

For Windows Server 2003, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2003 systems

- 1 Open the Control Panel (**Start>Control Panel**).
- 2 Double-click **Network Connections**, or right-click **Network Connections** and click **Open**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.

- Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 In the Public Status dialog box, in the General tab, click **Properties**.
- 6 In the Public Properties dialog box, in the General tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

For Windows Server 2008 systems, use the following procedure.

To verify the DNS settings and binding order for Windows Server 2008 systems

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, double-click **Manage network connections**.
- 4 Ensure the public network adapter is the first bound adapter:
 - From the Advanced menu in the Network Connections window, click **Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.

- Right-click the adapter for the public network and click **Status**.
- Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.

- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the **General** tab:
 - Select the **Internet Protocol Version 4 (TCP/IPv4)** check box.
 - Click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 12 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA 5.1 for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Veritas Cluster Server Database Agent for SQL. For a disaster recovery configuration, select the option to install GCO, and depending on your replication solution, select the option to install VVR or a hardware replication agent.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed by Microsoft.

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

[Table B-3](#) on page 1090 describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

Table B-3 Installation behavior with unsigned drivers

Driver Signing Setting	Installation behavior on the local system	Installation behavior on remote systems
Ignore	Always allowed	Always allowed
Warn	Warning message, user interaction required	Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation.
Block	Never allowed	Never allowed

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

Note: Windows Server 2008 does not enable you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 1091.

To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.
If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing Symantec Trusted certificate for unsigned drivers

The product installer provides an installation option for Symantec Trusted Software Publisher Certificate for Veritas Storage Foundation for Windows drivers that are not certified by Microsoft.

If you select this installation option, a Symantec Trusted certificate is installed on the systems you select for installation.

Warning: On Windows Server 2008, if this option is not selected, a remote install will not install any SFW drivers that are not certified by Microsoft. No notification is given and it will appear that installation was successful, but issues can arise later because the drivers were not installed.

If installing locally on Windows Server 2008, if this option is not selected, a driver signing popup will be displayed requiring user interaction.

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

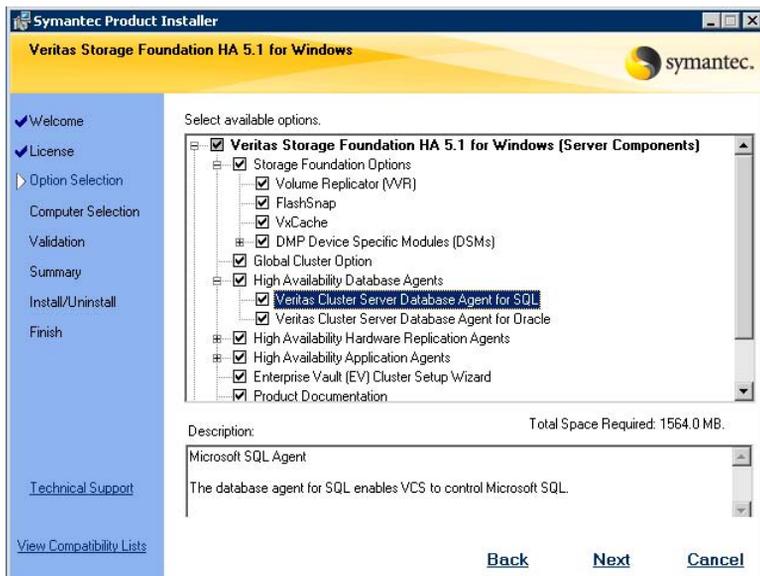
Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

To install the product

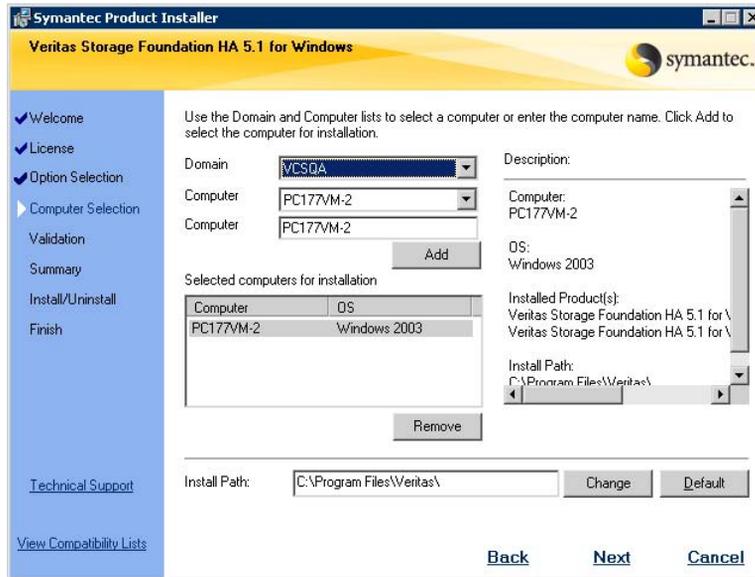
- 1 Allow the autorun feature to start the installation or double-click **Setup.exe**. The SFW Select Product screen appears.
- 2 Click **Storage Foundation HA 5.1 for Windows**.
- 3 Do one of the following:
 - Click **Complete/Custom** to begin installation.
 - Click the **Administrative Console** link to install only the client components.
- 4 Review the Welcome message and click **Next**.
- 5 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I accept the terms of the license agreement**, and then click **Next**.
- 6 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 7 Repeat for additional license keys. Click **Next**
 - To remove a license key, click the key to select it and click **Remove**.
 - To see the license key's details, click the key.

- 8 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:



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

9 Select the domain and the computers for the installation and click **Next**.



Domain

Select a domain from the list.

Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.

Computer

To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click **Add**.

To remove a computer after adding it, click the name in the Selected computers for installation field and click **Remove**.

Click a computer's name to see its description.

Install Path

Optionally, change the installation path.

- To change the path, select a computer in the Selected computers for installation field, type the new path, and click **Change**.
- To restore the default path, select a computer and click **Default**.

The default path is:

C:\Program Files\Veritas

For 64-bit installations, the default path is:

C:\Program Files (x86)\Veritas

- 10 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 11 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 12 If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning. At the Veritas Dynamic Multi-pathing warning, the time required to install the Veritas Dynamic Multi-pathing DSMs feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, connect only one physical path during installation. After installation, reconnect additional physical paths before rebooting the system.
- 13 Click **OK**.
- 14 Review the information and click **Install**. Click **Back** to make changes, if necessary.
- 15 The Installation Status screen displays status messages and the progress of the installation.
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 16 When the installation completes, review the summary screen and click **Next**.

- 17 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 18 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 19 Review the log files and click **Finish**.
- 20 Click **Yes** to reboot the local node.

Configuring VxSAS

Complete the following procedure to configure the VxSAS service for VVR.

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

Note: The VxSAS wizard will not be launched automatically after installing SFW or SFW HA. You must launch this wizard manually to complete the VVR security service configuration. For details on this required service, see *Veritas Storage Foundation Veritas Volume Replicator Administrator's Guide*.

To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxscfg.exe` from the command prompt of the required machine.
Read the information provided on the Welcome page and click **Next**.

2 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 the VxSAS service for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring the VxSAS service on the other hosts. Click **Next**.

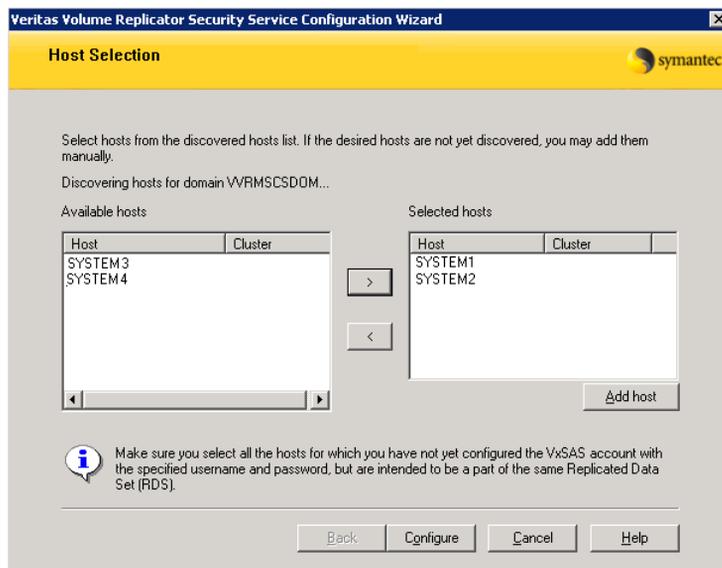
3 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

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 allows you to specify the domain name. Click **Add** to add the name to the Selected domains list.

Click **Next**.

4 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 name from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.

Adding a host

If the host name you require is not displayed, click **Add host**. In the Add Host dialog specify the required host name or IP in the **Host Name** field. Click **Add** to add the name to the Selected hosts list.

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring the VxSAS service.

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.
Click **Back** to change any information you had provided earlier.
- 6 Click **Finish** to exit the wizard.

Resetting the driver signing options

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

To reset the driver signing options

- 1 Open the Control Panel, and click **System**.
- 2 Click the **Hardware** tab and click **Driver Signing**.
- 3 In the Driver Signing Options dialog box, reset the option to **Warn** or **Block**.
- 4 Click **OK**.
- 5 Repeat for each computer.

Configuring the cluster

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

Complete the following tasks before creating a cluster:

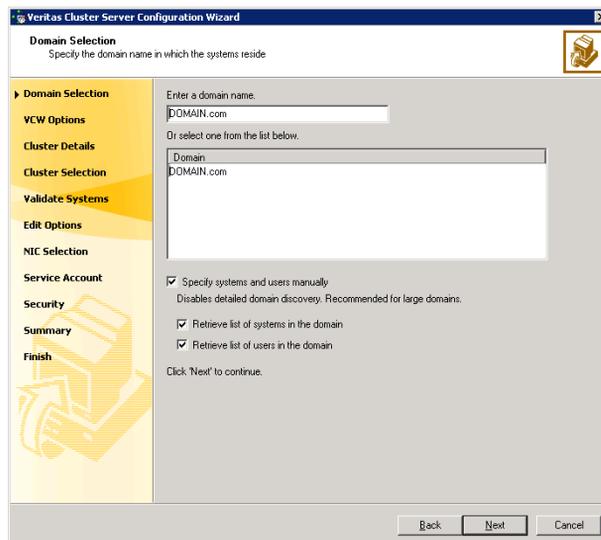
- Verify that each node uses static IP addresses (DHCP is not supported) and that name resolution is configured for each node.
- Set the required privileges:
 - You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
 - You must have administrative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

- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

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

To configure a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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

- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 8](#) on page 1102.

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

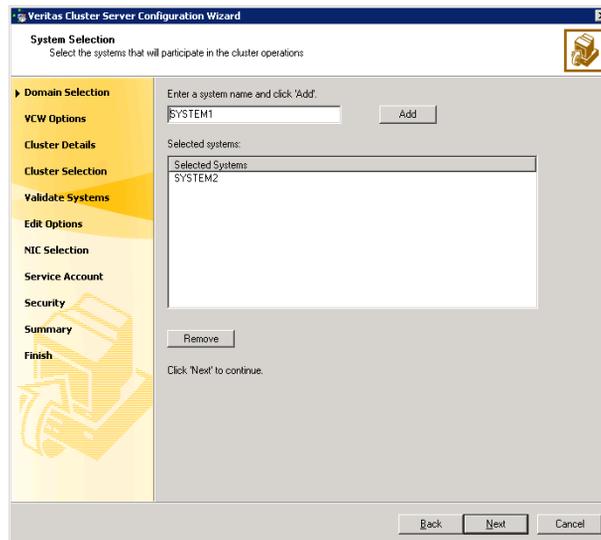
- Check the **Specify systems and users manually** check box.

Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.

■ Click **Next**.

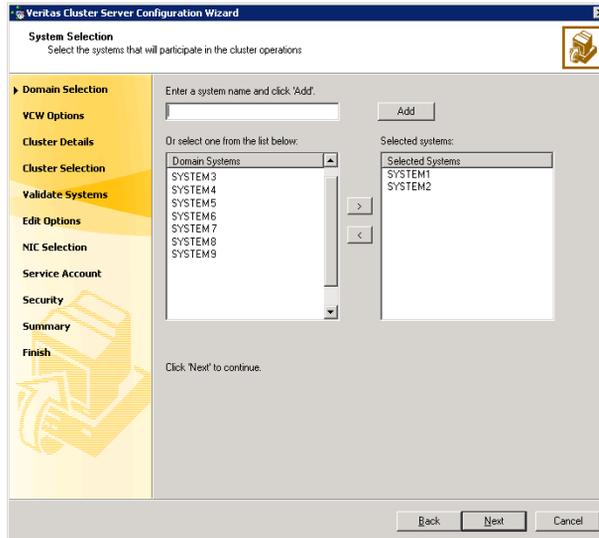
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 1102. Otherwise proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.



Proceed to [step 8](#) on page 1102.

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



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

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

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

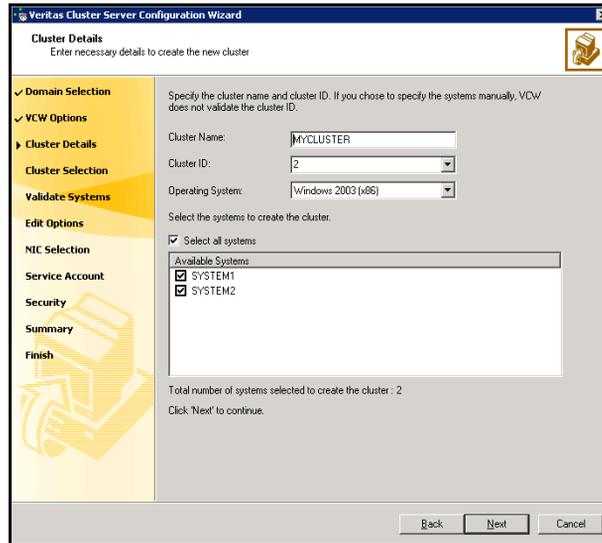
- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

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

Click **Next** to proceed.

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

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



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

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

Caution: If you chose to specify systems and users manually in [step 4](#) on page 1100 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 that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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. Check the **Select all systems** check box to select all the systems simultaneously.

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

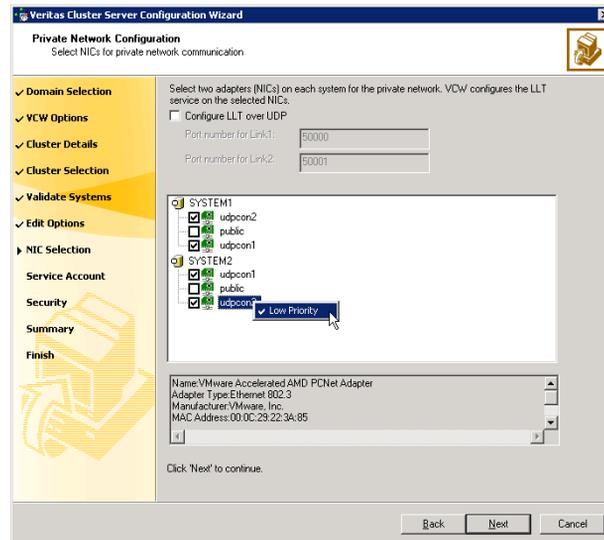
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.

If you chose to configure a private link heartbeat in [step 9](#) on page 1103, proceed to the next step. Otherwise, proceed to [step 12](#) on page 1106.

- 11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.

Do one of the following:

- To configure the VCS private network over Ethernet

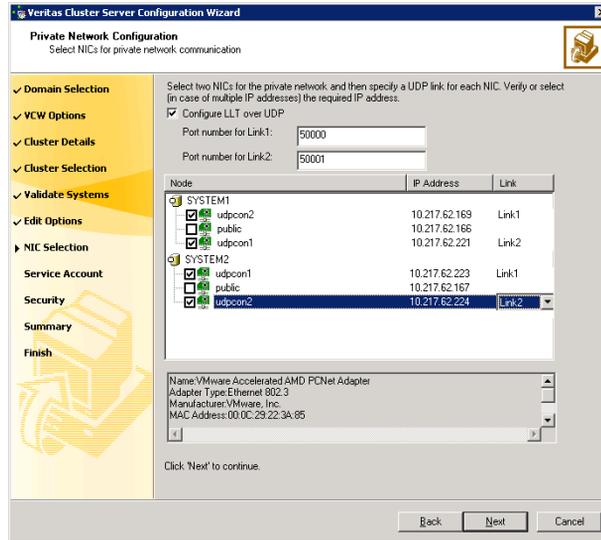


- Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC

address. Symantec recommends that you do not select teamed NICs for the private network.

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

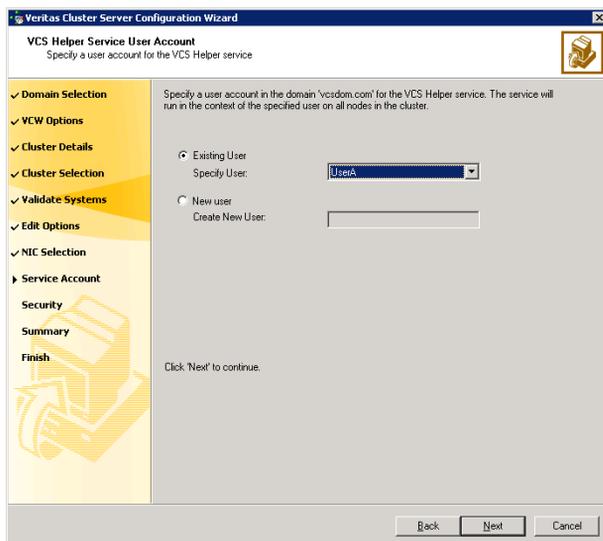
- To configure the VCS private network over the User Datagram Protocol (UDP) layer



- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
- Select the check boxes next to the two NICs to be assigned to the private network. Symantec 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. Each IP address can be in a different subnet.
 The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

- 12 On the VCS Helper Service User Account panel, specify the name of a domain user for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.



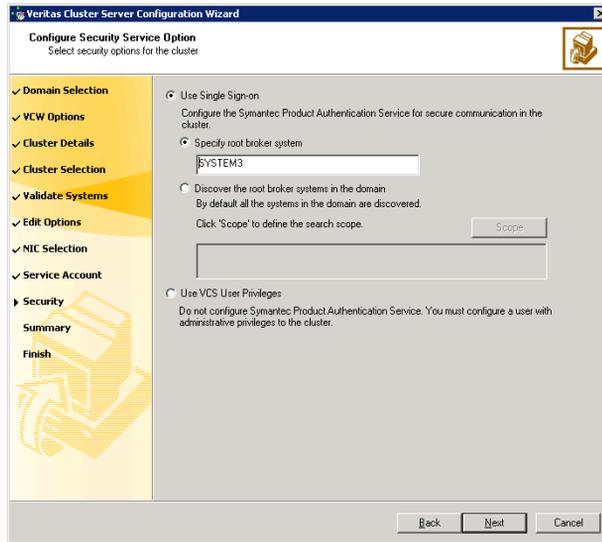
- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 1100, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the **Create New User** field, and then click **Next**.

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

- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

- 13 On the Configure Security Service Option panel, specify security options for the cluster and then click **Next**.
Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

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

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

For example, to search for systems managed by *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.

[Table B-4](#) on page 1108 contains some more examples of search criteria.

Table B-4 Search criteria examples

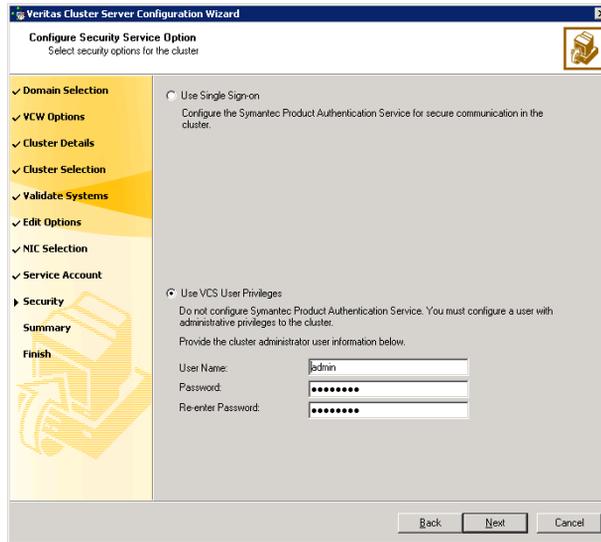
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .
Operating System	is (exactly)	*2003*	Displays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

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

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

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

- To use VCS user privilege



- Click **Use VCS User Privileges**.

The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.

Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 14 Review the summary information on the Summary panel, and click **Configure**.

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

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

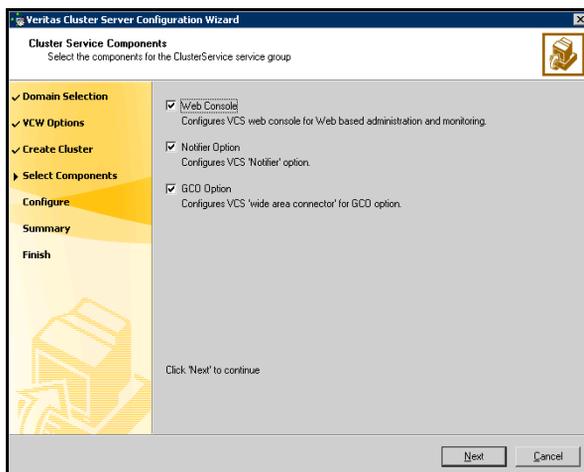
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Web Console, notification, and for global clusters.

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

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

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

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



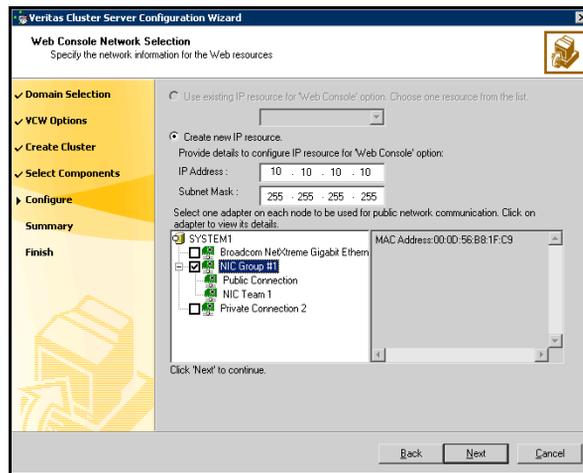
- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
The GCO Option applies only if you are configuring a Disaster Recovery environment and are not using the Disaster Recovery wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option

Configuring the Web Console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.



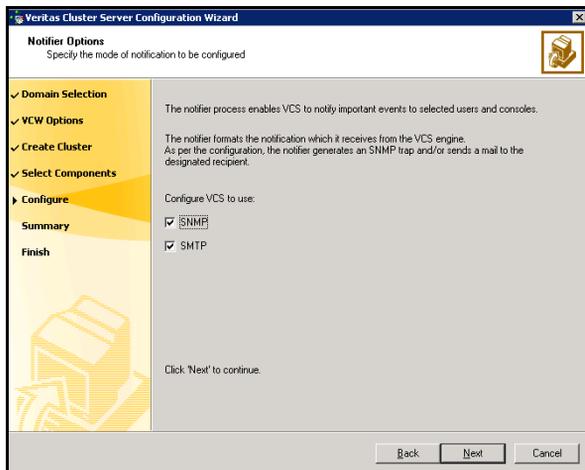
- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for 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 Web Console resources online when VCS is started, and click **Configure**.
 - 3 If you chose to configure a Notifier resource, proceed to [“Configuring notification”](#) on page 1112.
 If you chose to configure global cluster components, proceed to [“Configuring the wide-area connector process for global clusters”](#) on page 1115.
 Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

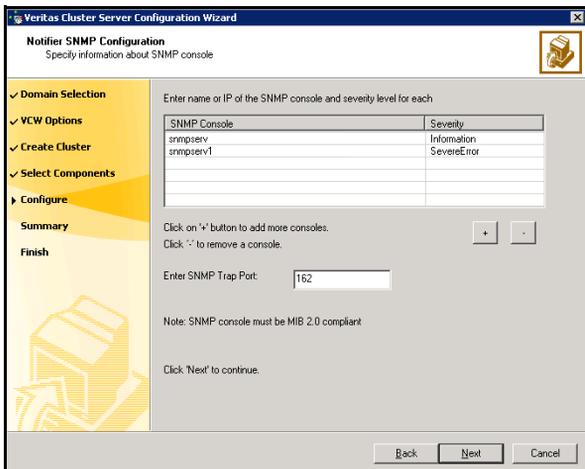
To configure notification

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

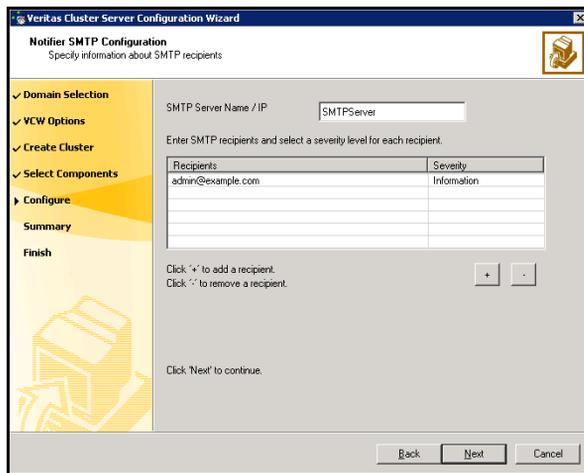


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

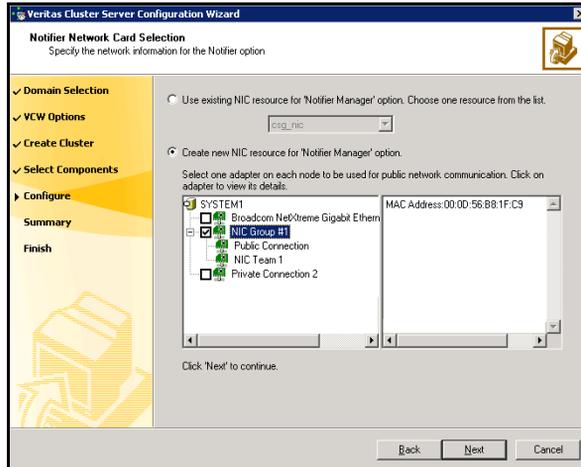


- Click a field in the SNMP Console column and enter the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click + to add a field; click - to remove a field.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



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

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



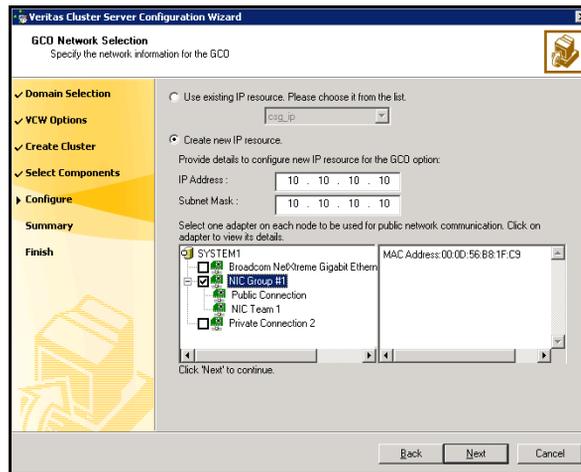
- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
 - If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started and click **Configure**.
 - 6 If you chose to configure global cluster components, proceed to [“Configuring the wide-area connector process for global clusters”](#) on page 1115. Otherwise, click **Finish** to exit the wizard.

Configuring the wide-area connector process for global clusters

This section describes steps to configure the wide-area connector resource required for global clusters.

To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
 - 3 Click **Finish** to exit the wizard.

Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your SQL Server database and logs.

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

A dynamic disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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 disk group from the current node and then importing it on the desired node.

Complete the following tasks before you create the disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place SQL Server use database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (INST1_DG) on shared disks and then create the following volumes:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

The following volumes may be created now or later in the configuration process.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files
- INST1_REPLOG: Contains the VVR Storage Replicator Log.

Caution: Do *not* assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption. VVR does not support the following types of volumes for the data and replicator log volumes: SFW (software) RAID 5 volumes, volumes with the Dirty Region Log (DRL) or Data Change Object (DCO), and volumes with commas in the names.

Use the following procedures to create the appropriate disk groups and volumes. Note that instructions for INST1_REPLOG are provided in “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 1151; you can wait until that chapter to configure VVR and its related volumes.

The first time you open the Veritas Enterprise Administrator (VEA), it will automatically connect to the localhost. It will also display the names of other hosts that are connected to VEA as nodes in the left side tree display.

Creating a cluster disk group

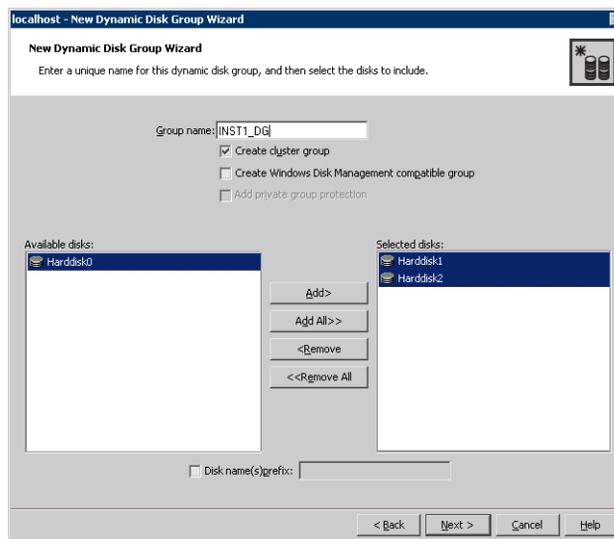
Create a cluster disk group on the first node of the cluster.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** 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.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Click the checkbox for **Create cluster group**.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

Note: For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

■ Click **Next**.

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

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.
- **INST1_DB1_VOL:** contains the user database files
- **INST1_DB1_LOG:** contains the user database log files
- **INST1_REPLOG:** Contains the VVR Storage Replicator Log.

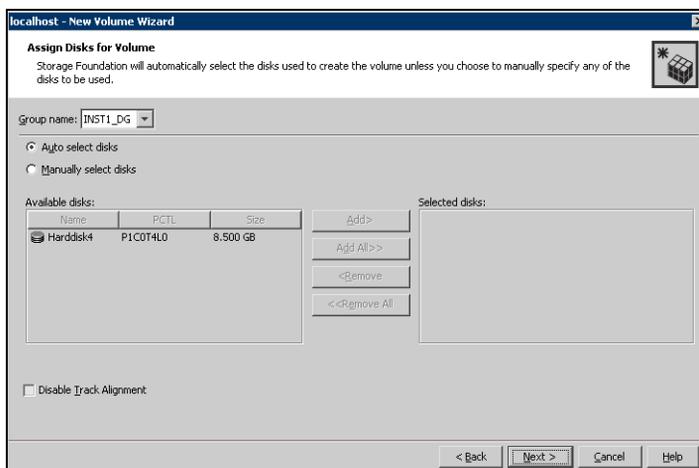
Caution: Do NOT assign a drive letter to the Replicator Log volume. This will limit access to that volume and avoid potential data corruption.

Refer to the steps below for the **INST1_DATA_FILES** and **INST1_REGREP_VOL** volumes. You can create the **INST1_DB1_VOL** and **INST1_DB1_LOG** volumes at this time or during the process of “[Creating a SQL Server user-defined database](#)” on page 1143. You can create the **INST1_REPLOG** volume at this time or during the process of “[Setting up the Replicated Data Sets \(RDS\)](#)” on page 1151.

Note: To ensure that the drive letters you assign to the new volumes will always be available on all nodes, assign drive letters starting in the middle of the alphabet. This way when drive letters are assigned as additional internal devices are added to a node there will not be a conflict with the volume drive letters.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** on the desktop. (Skip to step 4 if VEA is already connected to the appropriate host.)
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**. To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.



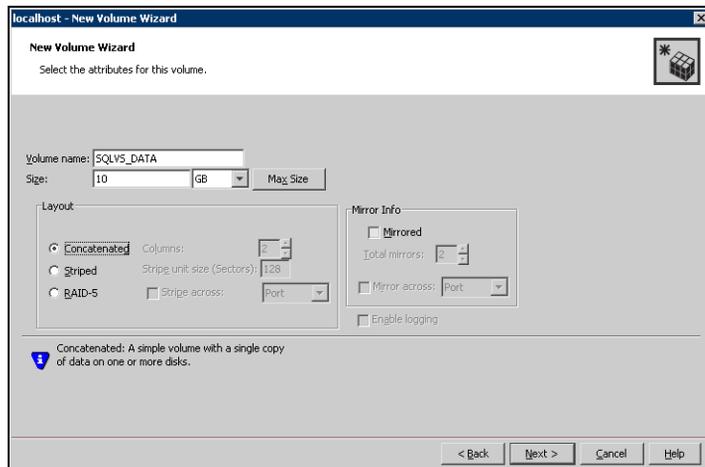
- Make sure the appropriate disk group name appears in the **Group name** drop-down list.

- Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

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

- Click **Next**.

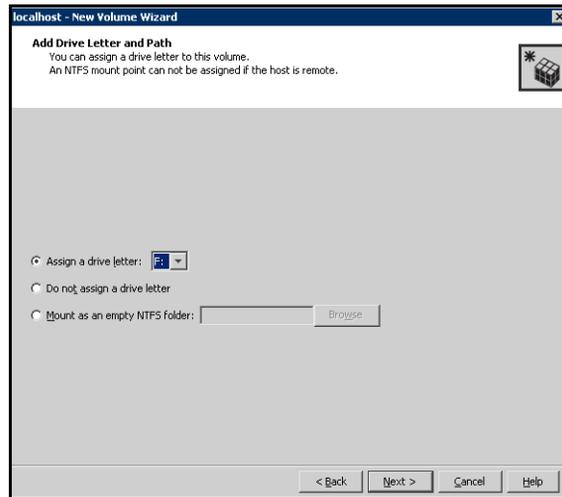
7 Specify the parameters of the volume.



- Enter a name for the volume. 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.
- 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.
- Verify that **Enable Logging** 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

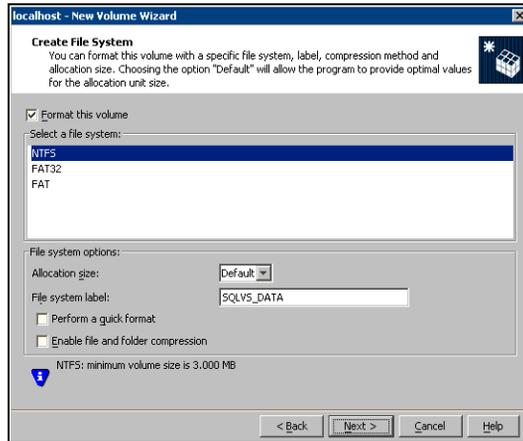
all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.



- 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.
- For the Replicator Log volume only:
Select **Do not assign a drive letter**.

9 Click **Next**.

10 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked.
 - For the Replicator Log volume only: Clear the Format this volume check box.
 - 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**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.

Installing and configuring SQL Server 2005 on the first node

In preparation for installing Microsoft SQL Server 2005, ensure the cluster disk group is imported to the first node and the volumes are mounted. Complete the following procedures to install and configure Microsoft SQL Server 2005:

- [Installing SQL Server 2005 on the first node](#)
- [Setting SQL Server 2005 services to manual start](#)

Installing SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

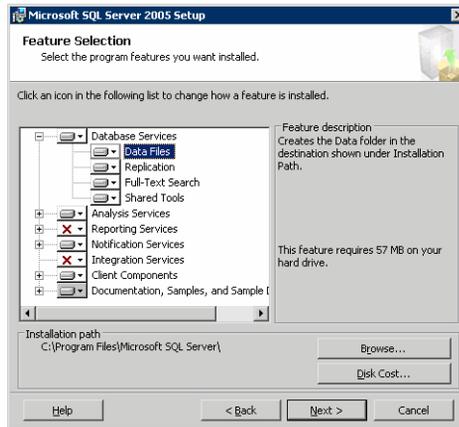
Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

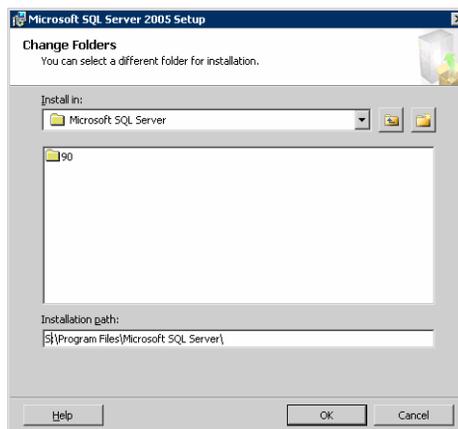
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.

- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.

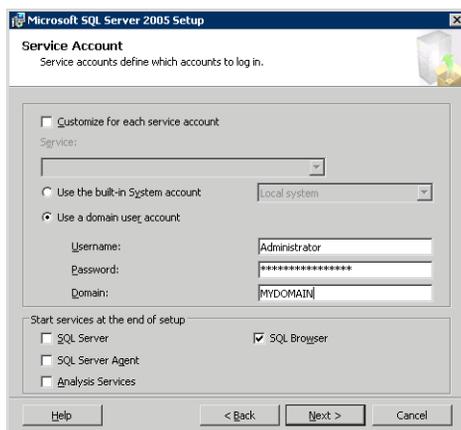


- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**. You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 1124, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.

- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:
 - Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.
See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.

Setting SQL Server 2005 services to manual start

Set all SQL Server 2005 services to manual start, except for the SQL Browser service. Ensure that the SQL Browser service is set to automatic.

To set the startup mode of SQL Server 2005 services

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance name and select **Properties**.
- 4 In the Properties dialog box, click the **Service** tab, select **Start Mode**, select **Manual** in the drop down list, and click **OK**.



- 5 Repeat for all other SQL Server services that are running on the server for this specific instance.

Preparing to install SQL Server 2005 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2005 Service”](#) on page 1128
- [“Deporting the cluster disk group”](#) on page 1129
- [“Importing the cluster disk group”](#) on page 1129
- [“Adding drive letters to mount the volumes”](#) on page 1130
- [“Renaming shared SQL Server 2005 files”](#) on page 1131

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 Service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Deporting the cluster disk group

In order to install SQL Server 2005 on additional nodes you must move ownership of the cluster disk group from the first node to an additional node. To move ownership, deport the clustered cluster disk group from the current node (SYSTEM1) and then import it to the desired node (SYSTEM2).

To deport the cluster disk group

- 1 Stop all processes accessing the volumes in the cluster disk group.
- 2 If necessary, click **Start > All Programs > Veritas > Veritas Enterprise Administrator** to start the Veritas Enterprise Administrator. If the Storage Foundation Assistant automatically opens, click **Close**.
- 3 Expand the host node and **Disk Groups** folder on the node where the cluster disk group is currently imported (SYSTEM1).
- 4 In the tree view, right-click the cluster disk group to be deported (INST1_DG) and select **Deport Dynamic Disk Group**.
- 5 Click **Yes** to deport the dynamic cluster disk group.

Importing the cluster disk group

To access a cluster disk group it must be imported to the node. Import the cluster disk group containing your SQL data files (INST1_DG) to the next node in the cluster (SYSTEM2).

To import a cluster disk group

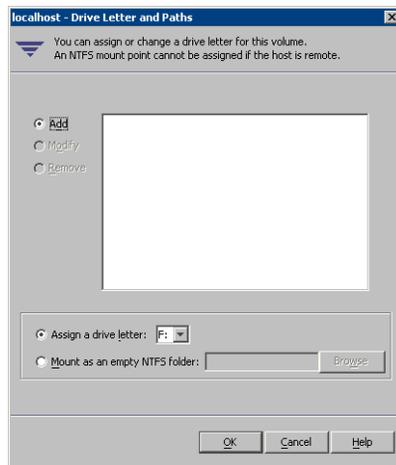
- 1 From the **Actions** menu select **Rescan** to update the disk information on the node where you want to import the cluster disk group.
- 2 Expand the host node and **Disk Groups** folder.
- 3 The cluster disk group will be visible on the node and will display the information (i) symbol.
 - In the tree view, right-click the cluster disk group and select **Import Dynamic Disk Group**.
 - Click **OK** in the **Import Dynamic Disk Group** dialog box.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the **Volumes** folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options 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 the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

Renaming shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server on the second node”](#) on page 1131
- [“Removing shared SQL Server files”](#) on page 1135

Installing SQL Server on the second node

Before installing Microsoft SQL Server 2005, verify that the cluster disk group is imported to the second node and the volumes are mounted (are assigned drive letters) See [“Importing the cluster disk group”](#) on page 1129 and [“Adding drive letters to mount the volumes”](#) on page 1130.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product.

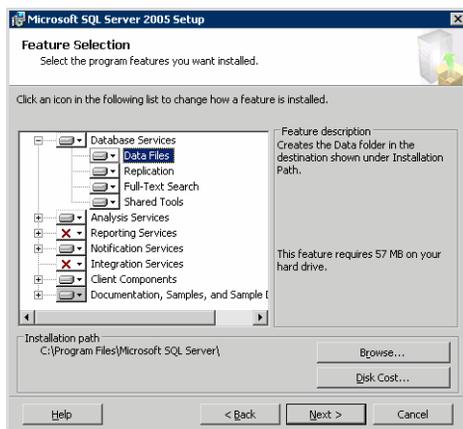
Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

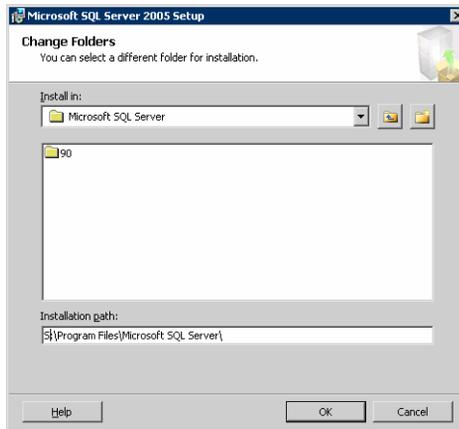
To install Microsoft SQL Server 2005

- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.
- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.
If you install optional components on one node, install the same components in the same order on other nodes.
- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:
 - Expand **Database Services**, click **Data Files**, and click **Browse**.



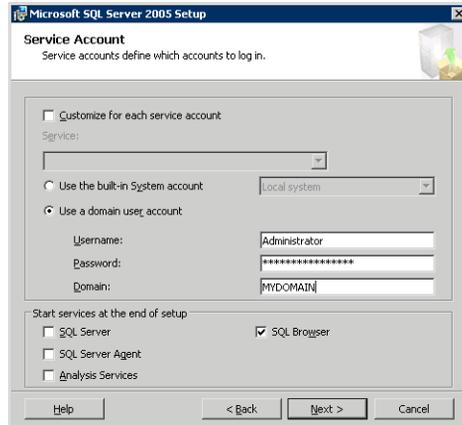
- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#) on page 1132, expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user name, password, and domain.

Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.

See Technote <http://support.veritas.com/docs/281828>.

- Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set all SQL services to manual start except for the SQL Browser service. Set the SQL Browser service to automatic.
- Refer to “[Setting SQL Server 2005 services to manual start](#)” on page 1127 for the procedure.

Repeat the procedures described in “[Preparing to install SQL Server 2005 on the second node](#)” on page 1128 and “[Installing SQL Server 2005 on the second node](#)” on page 1131 on any additional nodes.

Removing shared SQL Server files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, INST1-VS\INST1).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.

Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.

- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

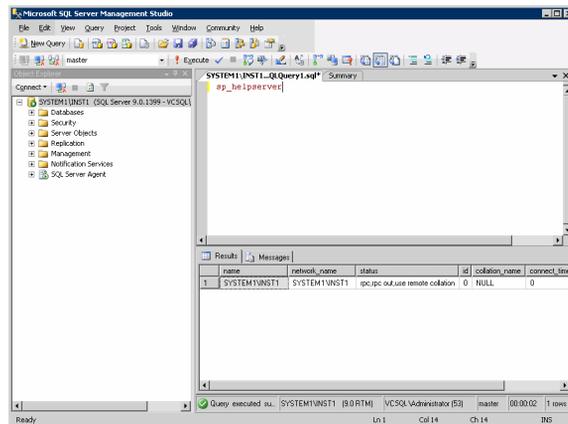
To set the internal name of the clustered instance

- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:
sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1 \ INST1. (For a default instance, you see only *System_Name* .)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:


```
sp_dropserver "System_Name\Instance_Name"
```

 where **System_Name\Instance_Name** is the name noted in [step 3](#) on page 1136.
 For example, for a named instance:
 sp_dropserver "SYSTEM1\INST1"
 For example, for a default instance:
 sp_dropserver "SYSTEM1"
 - Press F5.
 - 6 Delete the contents in the upper pane.

- 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
For example, for a named instance:
`sp_addserver "INST1-VS\INST1", local`
For example, for a default instance:
`sp_addserver "INST1-VS", local`
 - Press F5.
- 8 Exit the SQL Server Management Studio.
- 9 Stop the SQL instance on the node.
See "[Stopping the SQL Server 2005 Service](#)" on page 1128.

Configuring the VCS SQL Server service group

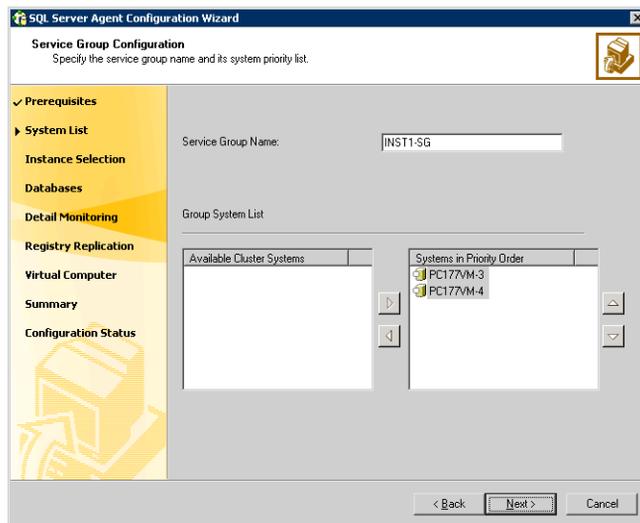
The Database Agent Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

- Verify that SFW HA, along with the VCS database agent for SQL Server, is installed on all cluster nodes. See "[Installing Veritas Storage Foundation HA for Windows](#)" on page 1090.
- Verify you have configured a VCS cluster using VCS Configuration Wizard (VCW). See "[Configuring the cluster](#)" on page 1099.
- Verify you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- Verify the drive containing the SQL Server 2005 system data files and registry replication information is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify the SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- Stop the SQL 2005 Server service for the SQL instance. See "[Stopping the SQL Server 2005 Service](#)" on page 1128.

- Assign a unique virtual server name to SQL Server 2005. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary site.
- Assign a unique virtual IP address to the SQL Server 2005 instance.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the

systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.

For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.

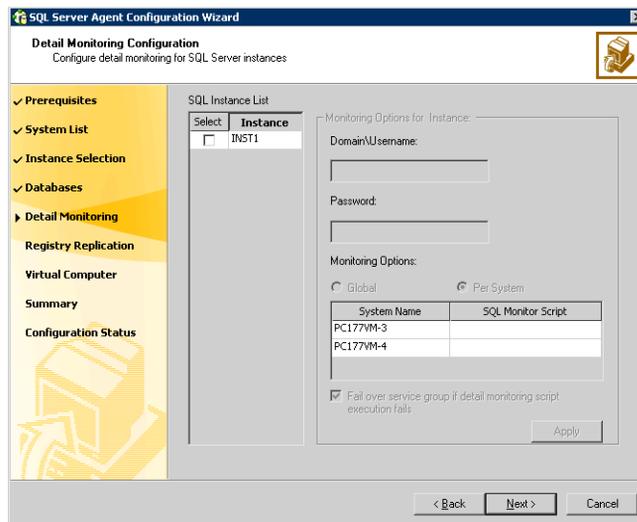
- Click **Next**.

- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment.

Click **Next**.

- 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

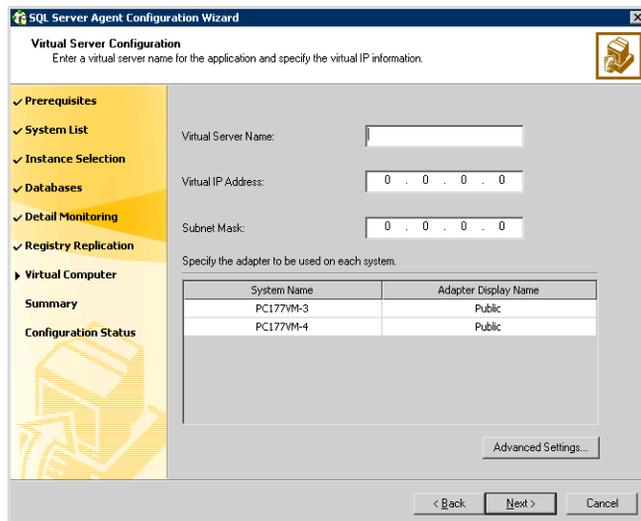
- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 1140 are available for selection.
- Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
- If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script

on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.

- Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) on page 1140 for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.
 - 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
 - 12 Configure the virtual server as follows:



- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.

- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
 - Enter the subnet mask to which the virtual IP address belongs.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
 - Click **Next**.
- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
- To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
- In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.

You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.
- The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.
- To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171.

Creating a SQL Server user-defined database

The following tasks enable you to use SFW HA to create and manage a SQL Server user-defined database.

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Configuration wizard to add the VMDg and MountV resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Refer to “[Creating volumes](#)” on page 1119 for information on how to use the VEA console to create a volume.

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Expand the icon associated with your server.
- 3 Right-click on **Databases** and select **New Database**.
- 4 In the New Database page, enter a name for the new database.
- 5 Click the browse button (...) in the **Path** column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 6 Select and edit other file properties as desired.
- 7 Click the browse button (...) in the **Path** column for the **Transaction Log** row and browse to the location of the volume you want to create for the transaction log, and click **OK**.

- 8 To add more data files if required:
 - a Select Add.
 - b Edit the properties in the new data file rows as required.
- 9 Click OK.

Adding VMDg and MountV resources

Before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted.

To add VMDg and MountV resources using the SQL Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.
- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the message informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:

- Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as **CRITICAL**.
- Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for disaster recovery”](#) on page 1171.

Verifying the cluster configuration

Failover simulation is an important part of configuration testing.

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

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

To switch service groups

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

To shut down an active cluster node

- 1 Gracefully shut down or restart the cluster node where the service group is online.
- 2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.
- 3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

- 4 If you need to move all the service groups back to the original node:
 - Restart the node you shut down in step 1.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.

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

Creating a parallel environment on the secondary site

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

Note: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

- [“Reviewing the prerequisites”](#) on page 1079
- [“Reviewing the configuration”](#) on page 1084
- [“Configuring the storage hardware and network”](#) on page 1087
- [“Configuring the cluster”](#) on page 1099
- [“Configuring cluster disk groups and volumes”](#) on page 1116

During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:

 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 1124

Select the same options at the secondary site as you did at the primary site.
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 1128

The instance name must be the same on the primary site and secondary site.
- [“Installing SQL Server 2005 on the second node”](#) on page 1131

- [“Configuring the VCS SQL Server service group”](#) on page 1138

Caution: Before you begin creating the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

The service group name and virtual computer name must be the same on both the primary site and secondary site.

Installing DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to [Appendix C, “Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 1149 after performing the tasks outlined in this chapter.

Configuring the DR components (VVR and GCO) without using the DR wizard

This appendix includes the following topics:

- [Reviewing the prerequisites](#)
- [Setting up the Replicated Data Sets \(RDS\)](#)
- [Creating the VVR RVG service group](#)
- [Configuring the global cluster option for wide-area failover](#)
- [Establishing secure communication within the global cluster \(optional\)](#)

This appendix provides information on configuring the following disaster recovery components: VVR, the Veritas Volume Replicator Agent for VCS, and the Global Cluster Option. Refer to the *Veritas Volume Replicator Administrator's Guide* for additional details on VVR.

You also have the choice of using array-based hardware replication for your disaster recovery solution. For information on configuring array-based hardware replication with VCS, see the VCS hardware agent documentation for the particular array you want to configure.

You configure these disaster recovery components for the primary and secondary sites after configuring the high availability and SQL components on both sites.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table C-1 Tasks for configuring the DR components

Objective	Tasks
“ Reviewing the prerequisites ” on page 1151	Verifying HA prerequisites for DR components
“ Setting up the Replicated Data Sets (RDS) ” on page 1151	Using the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites
“ Creating the VVR RVG service group ” on page 1162	Using the VVR Configuration Wizard to create a replication service group for the replicated volume group.
“ Configuring the global cluster option for wide-area failover ” on page 1164	<ul style="list-style-type: none"> ■ Linking clusters (adding a remote cluster to a local cluster) ■ Converting the application service group that is common to all the clusters to a global service group
“ Establishing secure communication within the global cluster (optional) ” on page 1169	Adding secure communication between local clusters within the global cluster (optional task)

Reviewing the prerequisites

The procedures assume that are deploying disaster recovery without using the DR wizard.

All tasks in [“Deploying disaster recovery: Manual implementation of a new SQL Server 2000 installation”](#) on page 1003, or [“Deploying disaster recovery: Manual implementation of a new SQL Server 2005 installation”](#) on page 1075 must be completed prior to starting this part of the DR solution.

Setting up the Replicated Data Sets (RDS)

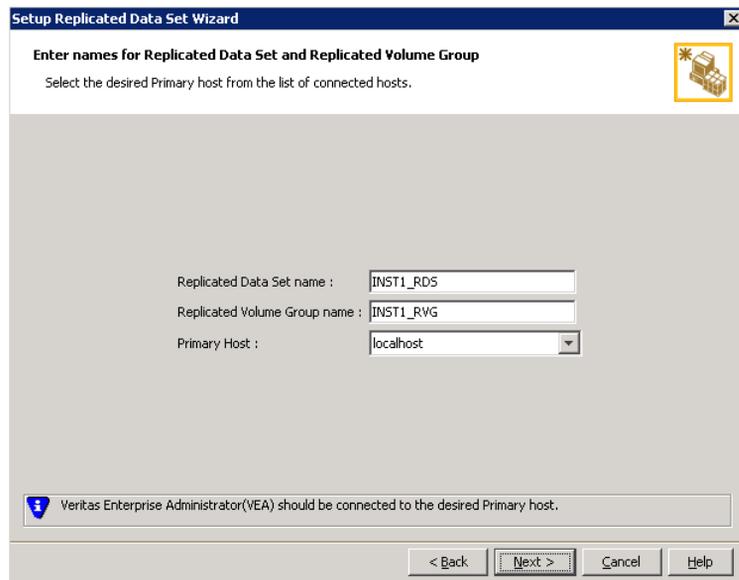
Set up the Replicated Data Sets (RDS) on the primary site and secondary site. You can configure an RDS using the Create RDS wizard for both sites.

- Verify that the data volumes are *not* of the following types as VVR does not support these types of volumes:
 - Storage Foundation for Windows (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
 - Volumes names containing a comma
- Verify that the cluster disk group is imported on the primary and secondary site

To create the Replicated Data Set

- 1 From the cluster node on the Primary where the cluster disk group is imported, launch the Veritas Enterprise Administrator (VEA):
 - Use the VEA console to launch the Setup Replicated Data Set Wizard.
OR
 - Launch the VEA by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**. From the VEA console, click **View > Connection > Replication Network**.
- 2 Right-click **Replication Network** and select **Set up Replicated Data Set**.
- 3 Read the Welcome page and click **Next**.

- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG).



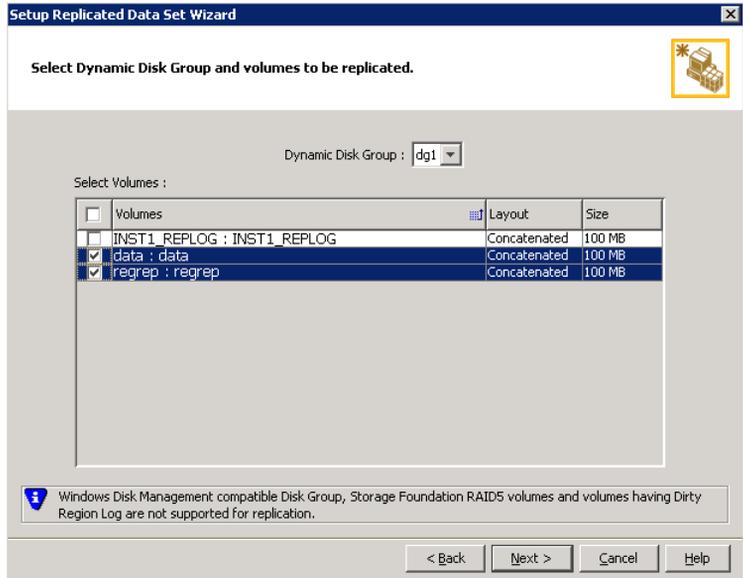
The screenshot shows a dialog box titled "Setup Replicated Data Set Wizard". The main heading is "Enter names for Replicated Data Set and Replicated Volume Group". Below this, there is a sub-heading: "Select the desired Primary host from the list of connected hosts." The dialog contains three input fields: "Replicated Data Set name" with the value "INST1_RDS", "Replicated Volume Group name" with the value "INST1_RVG", and "Primary Host" with a dropdown menu showing "localhost". At the bottom, there is a status bar with an information icon and the text "Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host." Below the status bar are four buttons: "< Back", "Next >", "Cancel", and "Help".

By default, the local host is selected as the **Primary Host**. To specify a different host name, make sure the required host is connected to the VEA console and select it in the **Primary Host** list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.

- 5 Click **Next**.

- 6 Select from the table the dynamic disk group and data volumes that will undergo replication.

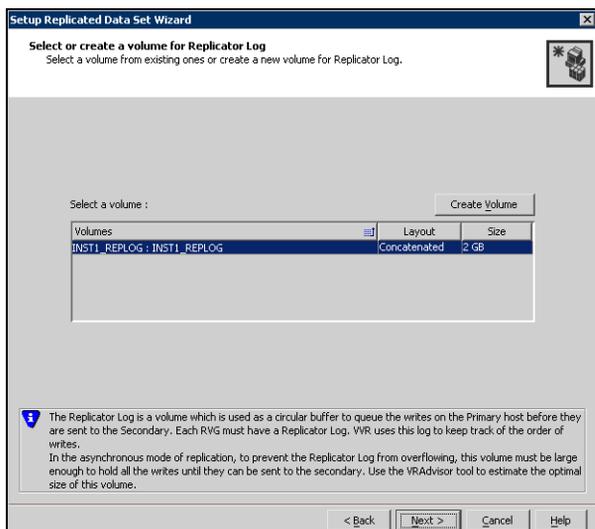


To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

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.

- 7 Click **Next**.

8 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 (INST1_REPLOG). If the volume does not appear in the table, click **Back** and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click **Create Volume** and enter the following information in the dialog box that displays.

- | | |
|-----------------------|--|
| Name | Enter the name for the volume in the Name field. |
| Size | Enter a size for the volume in the Size field. |
| Layout | Select the desired volume layout. |
| Disk Selection | <ul style="list-style-type: none">■ Choose Select disks automatically if you want VVR to select the disks for the Replicator Log.■ Choose Select disks manually to use specific disks from the available disks pane for creating the Replicator Log volume. Either double-click the disk to select it, or select Add to move the disks into the selected disks pane. |

- Click **OK** to create the Replicator Log volume.

- 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.

14 Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

The screenshot shows a window titled "Setup Replicated Data Set Wizard" with a sub-header "Edit replication settings". Below the sub-header is the instruction "Edit replication settings or click next." and a small icon of a server rack. The main area contains several configuration fields:

- Primary side IP: 10.217.53.214
- Secondary side IP: 10.217.53.215
- Replication Mode: Synchronous Override
- Replicator Log Protection: AutoDCM
- Primary RLINK Name: Pri_RLINK
- Secondary RLINK Name: Sec_RLINK

Below these fields is an "Advanced" button. At the bottom of the dialog, there is a warning message: "DHCP addresses are not supported by VVR." and four navigation buttons: "< Back", "Next >", "Cancel", and "Help".

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not

wish to modify basic properties then replication can be started with the default values when you click **Next**.

- | | |
|---------------------------|--|
| Primary side IP | Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address. |
| Secondary side IP | Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address. |
| Replication Mode | <p>Select the required mode of replication: Synchronous Override, Synchronous, or Asynchronous. The default is synchronous override.</p> <p>Synchronous Override enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous.</p> <p>Synchronous determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.</p> <p>Asynchronous determines updates from the application on the Primary site are completed after VVR updates in the Replicator Log. From there, VVR writes the data to the data volume and replicates the updates to the secondary site asynchronously.</p> <p>If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as MISSING.</p> |
| Replicator Log Protection | The AutoDCM is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows. |

The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.

The **Off** option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.

The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.

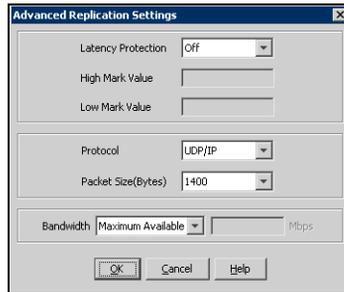
If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.

The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.

Primary RLINK Name	This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.
Secondary RLINK Name	This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then VVR assigns a default name.

Click **Next** to start replication with the default settings.

- Click **Advanced** to specify advanced replication settings. Edit the replication settings for a secondary host as needed.



Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

- **Off** is the default option and disables latency protection.
- **Fail** enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, VVR stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.
- **Override** enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Caution: Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.

Low Mark Value Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can “catch up” to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Protocol UDP/IP is the default protocol for replication.

Packet Size Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.

Bandwidth By default, VVR uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.

Click **OK** to close the dialog box.

16 Click **Next**.

17 On the **Start Replication** page, select **Start Replication**.

Synchronize
Automatically

If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Veritas Storage Foundation™ Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu. Click **Next** to display the Summary page.
- 18 Review the information.
Click **Back** to change any information you had specified and click **Finish** to add the secondary host to the RDS and exit the wizard.

Creating the VVR RVG service group

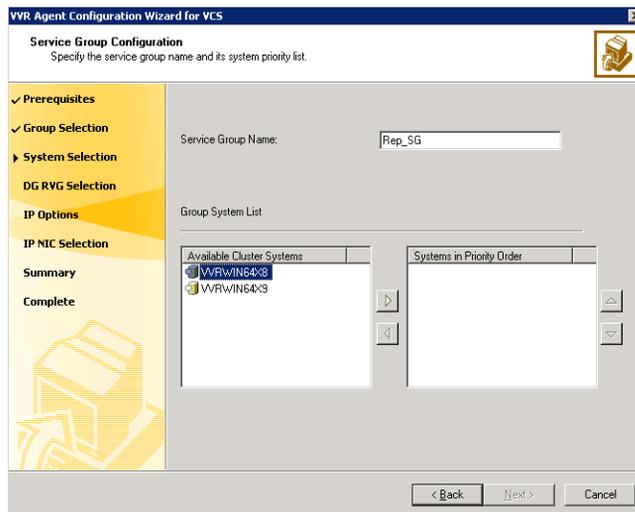
Prerequisites:

- Verify that the disk group is imported on the node on which you want to create the Replication Service Group.
- Verify VCS is running, by running the following command on the host on which the you intend to run the Volume Replicator Agent Configuration Wizard.

```
> hasys -state
```

To create a replication service group

- 1 From the active node of the cluster at the primary site, click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard** to launch the configuration wizard.
- 2 Review the requirements on the Welcome page and click **Next**.
- 3 In the **Wizard Options** panel, click **Create a new replication service group** and click **Next**.
- 4 Specify the service group name and system priority list as follows:



- Enter the service group name (INST1_RVG_GRP).
- In the **Available Cluster Systems** box, click the nodes on which to configure the service group, and click the right-arrow icon to move the nodes to the service group's system list. Make sure that the set of nodes

selected for the replication service group is the same or a superset of nodes selected for the SQL Server service group. Ensure that the nodes are in the same priority order.

- To remove a node from the service group's system list, click the node in the **Systems in Priority Order** box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the **Systems in Priority Order** box, then click the up and down arrow icons. The node at the top of the list has the highest priority.
 - Click **Next**.
- 5 A message appears, indicating that the configuration will be changed from Read Only to Read/Write. Click **Yes** to continue.
 - 6 In the Disk Group and Replicated Volume Group Configuration panel, make the following selections:
 - Select **Configure RVGPrimary resource for selected RVG**.
 This resource is required when you want to configure your setup to automatically enable takeover in case of a failure of the Primary cluster. The RVGPrimary resource is created in the application service group and replaces the VMDg resource.
 - Select the replicated volume group for which you want to configure the RVG primary resource.
 - Click **Next**.
 - 7 In the IP Resource Options panel, select **Create a new IP resource** and click **Next**.
 - 8 In the Network Configuration panel, enter the network information as follows:
 - Verify or enter the virtual IP address; use the IP address specified as the primary IP address when you configured the RDS.
 - Specify the subnet mask.
 - Specify the adapters for each system in the configuration.
 - Click **Next**.
 - 9 Review the summary of the service group configuration as follows:

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

 - If necessary, change the resource names; the wizard assigns unique names to resources based on their respective name rules.
 - To edit a resource name, click the resource name and modify it. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.

- Click **Next** to create the replication service group.
- 10 A warning informing you that the service group will be created is displayed. When prompted, click **Yes** to create the service group.
- 11 Click **Finish** to bring the replication service group online.
- 12 Check the prerequisites, then repeat the wizard at the secondary site, specifying the appropriate values.
The name for the application service group must be the same on both sites. The SQL service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

Configuring the global cluster option for wide-area failover

The Global Cluster option is required to manage global clustering for wide-area disaster recovery. Creating a global cluster environment involves:

- Connecting standalone clusters by adding a remote cluster to a local cluster.
- Converting the local service group that is common to all the clusters to a global service group.

Use the VCS Java Console or the Cluster Management Console (Single Cluster Mode) also referred to as Web Console, to perform global cluster operations; this guide only provides procedures for the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on GCO operations available from the Java and Web Consoles.

Prerequisites

Creating a global cluster environment requires:

- All service groups properly configured and able to come online.
- The service group serving as the global group has the same unique name across all applicable clusters.
- The clusters use the same version of VCS.
- The clusters use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.
- The names of the clusters at the primary and secondary sites and the virtual IP addresses associated with them are registered in the DNS with reverse lookup.

Linking clusters: Adding a remote cluster to a local cluster

The VCS Java Console provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

- 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 Cluster Explorer:

- 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.

Note: Symantec does not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

To add a remote cluster in Cluster Explorer

- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
- 2 Review the required information for the **Remote Cluster Configuration Wizard** and 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 as follows:
If the cluster is not running in secure mode:

- 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.
- Enter the password.
- Click **Next**.

If the cluster is running in secure mode:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.
- If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- Click **Next**.

- 5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.
- 6 Verify that the heartbeat connection between clusters is alive. From the command window enter `hahb -display`. The state attribute in the output should show **alive**.
If the state is **unknown**, then offline and online the ClusterService group.

Converting a local service group to a global service group

After linking the clusters, use the Global Group Configuration wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

To convert a local service group to a global group

- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster, and click **Configure Global Groups**.
or

From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3b.

- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify:
 - 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:

- | | |
|----------------------------|---|
| Cluster not in secure mode | <ul style="list-style-type: none"> ■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system. ■ Verify the port number. ■ Enter the user name. ■ Enter the password. ■ Click OK. ■ Repeat these steps for each cluster in the global environment. |
|----------------------------|---|

- Cluster in secure mode
- 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

After converting the local service group that is common to the global clusters to a global group, use the Cluster Explorer to bring the global service group online.

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:
 - Click the remote cluster to bring the group online.
 - Click the specific system, or click **Any System**, to bring the group online.
 - Click **OK**.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

To add the `-secure` option to the `StartProgram` resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add `-secure` switch to the path of the executable Scalar Value.
For example:

```
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe"  
-secure
```
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.

Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel  
<low|medium|high> [--hashfile <filename> | --hash <root  
hash in hex>]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:

from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

from RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the `ClusterService-Proc (wac)` resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the **Process** agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

Configuring an MSDTC service group for disaster recovery

This appendix includes the following topics:

- [Tasks for configuring MSDTC for disaster recovery](#)
- [Reviewing the prerequisites](#)
- [Reviewing the configuration](#)
- [Configuring cluster disk groups and volumes](#)
- [Mounting drives used by the MSDTC service group](#)
- [Creating an MSDTC service group](#)
- [Creating an MSDTC client](#)
- [Setting up the secondary site: Creating a parallel environment](#)
- [Installing DR components on the primary and secondary sites](#)

Tasks for configuring MSDTC for disaster recovery

The VCS Database Agent for SQL, which you install during SFW HA installation for Microsoft SQL high availability, includes an MSDTC agent. The MSDTC agent can provide high availability for the Microsoft Data Transaction Coordinator (MSDTC) service. The MSDTC agent comprises two parts: MSDTC client and MSDTC server.

You can configure high availability for MSDTC either before or after configuring high availability for Microsoft SQL. To configure high availability for MSDTC,

you use the SQL Server Configuration Wizard to configure an MSDTC service group for the MSDTC server and to configure the MSDTC client.

To modify an existing MSDTC service group, see the *Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide*.

After configuring high availability for MDSTC on the primary site, you can configure MSDTC on the secondary site for disaster recovery.

The table below outlines the high-level objectives and the tasks to complete each objective:

Table D-1 Tasks for configuring MSDTC for disaster recovery

Objective	Tasks
“Reviewing the prerequisites” on page 1173	Verifying hardware and software prerequisites
“Reviewing the configuration” on page 1173	<ul style="list-style-type: none"> ■ Reviewing MSDTC service group configuration ■ Reviewing the sample configuration
“Mounting drives used by the MSDTC service group” on page 1181	Mounting the drives used by the service group
“Configuring cluster disk groups and volumes” on page 1177	<ul style="list-style-type: none"> ■ Configuring disk groups for an MSDTC service group ■ Configuring volumes for an MSDTC service group
“Creating an MSDTC service group” on page 1182	Creating an MSDTC service group
“Creating an MSDTC client” on page 1184	Creating an MSDTC client
“Setting up the secondary site: Creating a parallel environment” on page 1185	<ul style="list-style-type: none"> ■ Reviewing the prerequisites ■ Reviewing the MSDTC service group configuration ■ Mounting the drives used by the service group ■ Configuring disk groups for an MSDTC service group ■ Configuring volumes for an MSDTC service group ■ Creating an MSDTC service group ■ Creating an MSDTC client

Reviewing the prerequisites

You must meet the following prerequisites before creating and configuring the MSDTC service group:

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that the VCS Database Agent for SQL is installed on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Configuration Wizard (VCW).
- Verify that the drives containing the MSDTC logs and registry replication directory are mounted on one node (the node on which you are configuring the service group) and unmounted on all other nodes.
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- Assign a unique virtual server name and virtual IP address for the MSDTC server.
- Verify that the Distributed Transaction Coordinator service is stopped.

Reviewing the configuration

MSDTC servers can coexist with SQL servers on the same cluster nodes. If the MSDTC server and the SQL server are running on the same node, the MSDTC client is left in the default configuration. If the MSDTC server is not configured on the same node as the SQL server, then the MSDTC client must be configured on that node using the SQL Configuration Wizard.

For instance, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage. The shared storage can be managed using Veritas Storage Foundation for Windows (SFV). Two configurations are possible:

- SQL server is configured on different nodes than the MSDTC server
- SQL server is configured on the same node as the MSDTC server

Figure D-1 MSDTC server configured on different nodes than SQL server (primary site)

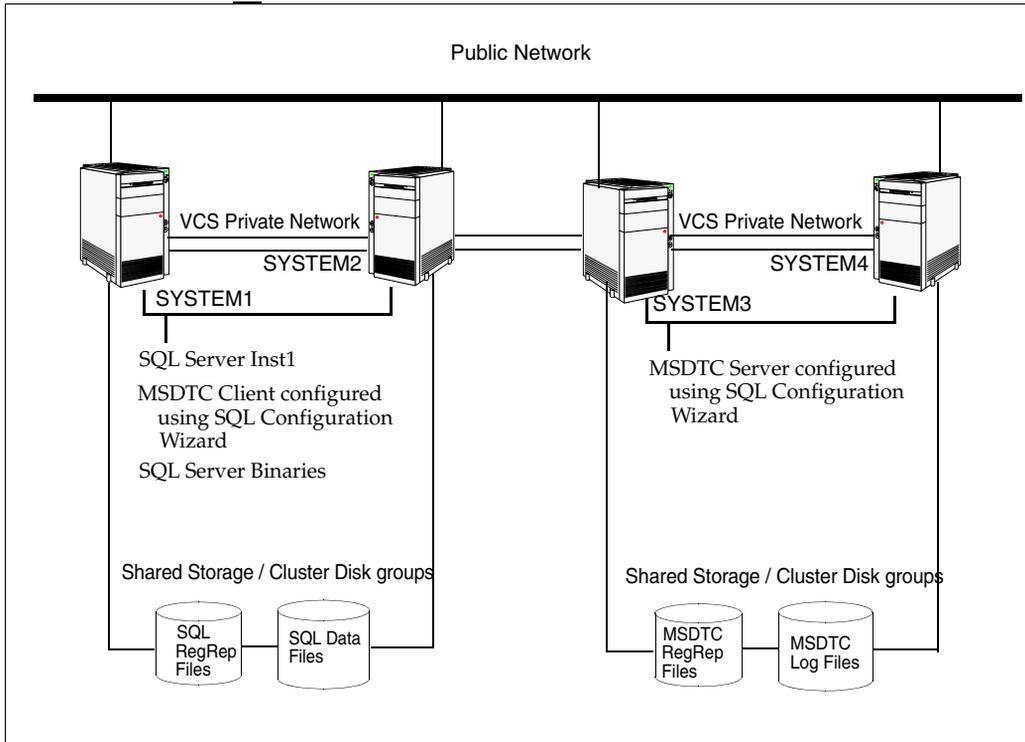
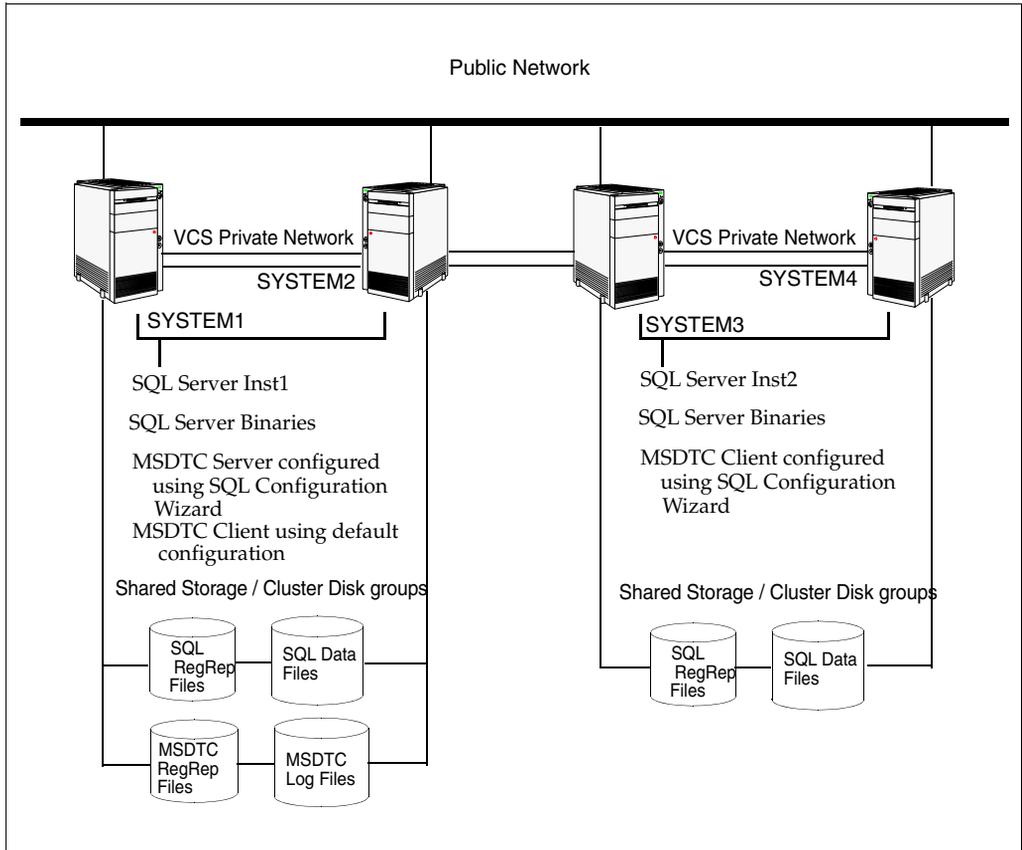
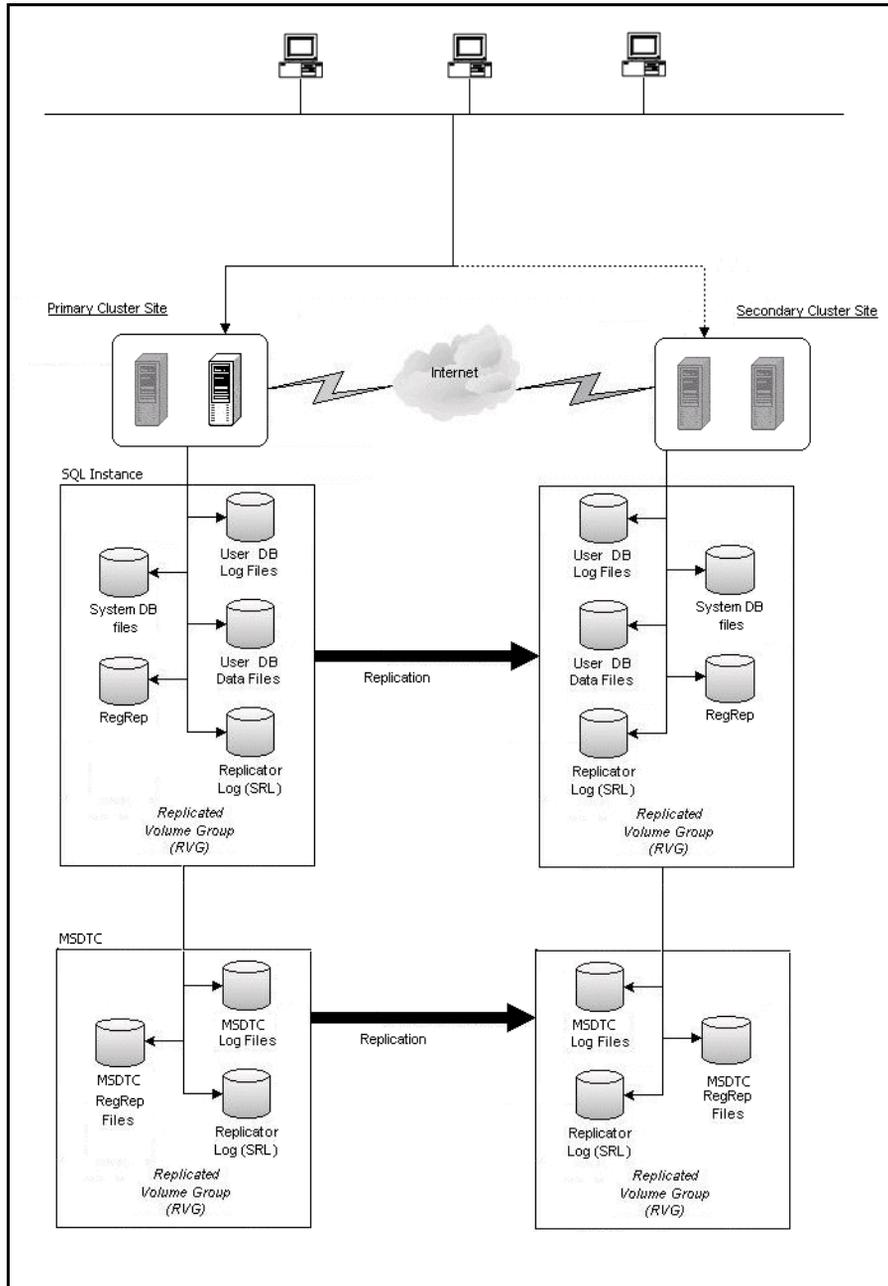


Figure D-2 MSDTC server configured on the same node as SQL server (primary site)



For replication, the secondary site must be configured exactly the same as the primary site.

Figure D-3 Replication configuration for MSDTC service group



Configuring cluster disk groups and volumes

Create a cluster disk group and volumes to manage your MSDTC service group.

A disk group is a collection of disks that is imported or deported as a single unit. SFW uses disk groups to organize disks or LUNs for management purposes. 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: 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: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (*VEA - Control Panel - System Settings*). See the Veritas Storage Foundation Administrator's Guide for more information.

Complete the following tasks before you create the cluster disk group and volumes:

- Determine the layout or configuration for each volume and the total number of disks needed. Symantec recommends that you place MSDTC files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the **Expand Volume** command but you can not decrease the size.
- Verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must **Rescan**, and if necessary, use the **Write Signature** command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

On the first node of the cluster you will first need to create a cluster disk group (MSDTC_DG) on shared disks and then create the following volumes:

- MSDTC_LOG: contains the MSDTC log files.
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB volume for this purpose.

Creating a cluster disk group

Create a cluster disk group on the first node of the cluster.

To create a cluster disk group

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 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 On the **Welcome** page of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group you want to create.
 - Enter the name of the disk group (for example, MSDTC_DG).
 - Click **Create cluster group**.
 - Select the appropriate disks in the **Available disks** list, and click the **Add** button to move them to the **Selected disks** list.
 - Click **Next**.
- 7 Review the selected disks and click **Next**.
- 8 Review the summary information and click **Finish**.

Creating volumes

This section will guide you through the process of creating a volume on a cluster disk group. Repeat the procedure below to create the following volumes on the first node of the cluster:

- MSDTC_LOG: contains the MSDTC log files
- MSDTC_REGREP: contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group.

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

To create a volume

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the 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 the disk group you created for the MSDTC volumes and select **New Volume**.
- 5 On the **Welcome** page of the New Volume Wizard, click **Next**.
- 6 On the **Assign Disks for Volume** page, make sure the name of the disk group that you created for MSDTC appears in the Group name field.
- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended. After selecting the necessary disks, click **Next**.
- 8 Specify the volume attributes:
 - Enter a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - 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.
- 9 In the **Add Drive Letter and Path** page, 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.
- 10 Specify information about the attributes for the file system that will be created:
- 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**.
- 11 Review the Summary dialog box and click **Finish**.
- 12 Repeat the procedure as necessary to create the additional volumes needed for your deployment. Create the volumes on the first node of the cluster only.

Mounting drives used by the MSDTC service group

Ensure that the shared volumes created for the MSDTC log files and registry replication information are mounted on the node where you run the SQL Server Configuration Wizard to configure the service group. The volumes must be unmounted on all other nodes in the cluster.

If you created the MSDTC disk group and volumes on another node, mounting a volume involves deporting the disk group from the other node and then importing the disk group to the node where you are configuring the service group.

Occasionally, when a disk group is imported, a drive letter may not be associated with an existing volume. If this occurs, use VEA to add a drive letter or folder path and mount the volume so that it can be seen by the operating system.

To mount volumes

- 1 If the disk group is not imported, deport it from the other node and then import it on the current node:
 - Stop all processes accessing the volumes in the disk group, and from the VEA console's tree view, right-click the disk group and select **Deport Dynamic Disk Group**. Then select **Yes**.
 - In the VEA, connect to the node where you want to import the cluster disk group. In the tree, right-click **Storage Agent**, and click **Rescan** to update the disk information on the node. Then right-click the disk group and select **Import Dynamic Group** and click **OK**.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.
- 3 In the Drive Letter and Paths dialog box, select **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder. Assign the same drive letter or mount path that was assigned when the volume was created
 - *To assign a drive letter*
In the Assign Drive Letter dialog box, select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
 - *To mount the volume as a folder*
Select the **Mount as an empty NTFS folder option** and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.
- 6 Repeat [step 2](#) on page 1181 through [step 5](#) on page 1181 for all the volumes to be mounted. You can now create the MSDTC service group.

Creating an MSDTC service group

MSDTC is a global resource and is accessed by more than one SQL Server service group. Symantec recommends you to configure one MSDTC service group in a VCS cluster.

Note: You must be a Local Administrator on the node where you run the wizard.

To configure an MSDTC service group

- 1 If you have just configured a SQL service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 2 Select **MSDTC Server - Service Group Configuration** and **Create**. Click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the service group name and system list.
 - Enter a name for MSDTC service group.
 - In the **Available Cluster Systems** box, select the systems on which to configure the service group and click the right-arrow to move the systems to the service group's system list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - To change a system's priority in the **Systems in Priority Order** list, select the system and click the up and down arrows. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 - Click **Next**. If the configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 Specify the information related to the virtual server.
 - Enter a virtual server name for the node on which the DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.
 - Enter a unique virtual IP address for the MSDTC server.
 - Enter the subnet mask to which the virtual IP address belongs.

- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system.
 The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop down list.
 - Click **Next**.
- 6 Specify the MSDTC log and replication directory and click **Next**. Symantec recommends using different paths for these directories. If the directory does not exist, the wizard creates it.
 - 7 Review the service group configuration, change the resource names if desired and click **Next** to create the service group.
 - The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.
 - The wizard assigns unique names to resources. Change names of the resources, if desired. To edit a resource name, select the resource name and either click it or press the **F2** key. Press **Enter** after editing each resource name. To cancel editing a resource name, press **Esc**.
 - 8 A message appears informing you that the wizard will run commands to modify the service group configuration. Click **Yes** to create the service group.
 Various messages indicate the status of these commands.
 - 9 In the Configuration Complete panel, check **Bring the service group online** to bring the configured service group online. To bring the service group online later, uncheck the option.
 - 10 Click **Next** to create an MSDTC client, or click **Finish** to exit the wizard.

Creating an MSDTC client

Set the MSDTC client to run on nodes where:

- A SQL instance is configured to run
- The MSDTC server is not configured to run.

Before configuring the MSDTC configuration wizard to configure an MSDTC client:

- Verify the MSDTC service group is online in the cluster.
- Verify the node on which you run the wizard is not a part of an MSDTC service group system list.

To configure an MSDTC client

- 1 If you have just configured the MSDTC service group and you are in the Configuration Wizard, proceed to the next step. Otherwise, start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 2 In the Select Configuration Option panel, select **MSDTC Client Configuration** and click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 In the System Selection panel, specify the nodes on which the MSDTC client will be configured and click **Next**:
 - Select the nodes in the **Available Cluster Systems** list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - Click the right arrow to add them to the **Selected systems** list.
 - The **Unavailable Cluster Systems** lists the nodes that have an MSDTC service group configured and are therefore not available for setting up an MSDTC client.
- 5 If the MSDTC service group is not online in the cluster, an informational message appears informing you that the wizard will bring the MSDTC service group online. Click **Yes**.
- 6 On the Specify DTC Node panel, specify the virtual DTC server name and click **Next**.
- 7 On the Configuration Complete panel, click **Finish** to exit the wizard.

Setting up the secondary site: Creating a parallel environment

After setting up a SFW HA environment on the primary site, use the guidelines in this chapter to complete the same tasks on the secondary site:

- [“Reviewing the prerequisites”](#) on page 1173
- [“Reviewing the configuration”](#) on page 1173
- [“Configuring cluster disk groups and volumes”](#) on page 1177
During the creation of disk groups and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Cluster disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- [“Creating an MSDTC service group”](#) on page 1182
Select the same options at the secondary site as you did at the primary site.
- [“Creating an MSDTC client”](#) on page 1184
The instance name must be the same on the primary site and secondary site.

Caution: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, and the SQL virtual IP resource on the primary site. The remaining resources should be online, including the VMDg resource.

The service group name, virtual computer name, and IP address must be the same on both the primary site and secondary site.

Installing DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to [“Configuring the DR components \(VVR and GCO\) without using the DR wizard”](#) on page 1149 after performing the tasks outlined in this chapter. You can configure SQL Server first if you have not already configured it, and then proceed to the DR configuration tasks.

Index

A

- active/active configuration
 - illustration 60, 210
 - IP addresses 62, 212
 - key information 66, 216
 - overview 60
 - service group requirements 119, 272
- active/passive configuration
 - illustration 58, 208
 - IP addresses 59
 - overview 58, 208

C

- campus cluster
 - configuration overview 388, 462
 - defined 372
 - disk space requirements 383, 457
 - failover using the forceimport attribute 374
 - forceimport 449, 524
 - new installation 379, 453
 - overview 371, 374
 - preconditions for a cluster disk group 416, 490
 - process overview 379, 453
 - site failure 449, 524
 - SQL Server service group, modifying the IP resource 448, 523
 - verifying cluster configuration 449, 524
- cloning for DR
 - secondary storage (array-based replication) 833, 935
 - secondary storage (VVR replication) 829, 931
 - service group for SQL Server 2000 850
 - service group for SQL Server 2005 954
- Cluster Management Console 127
- clusters
 - assigning user privileges 823, 926
 - configuring the cluster 86, 237, 558, 676, 799, 902, 1024, 1097
 - configuring the hardware and network 68, 218, 547, 664, 789, 892, 1012, 1085
 - setting the internal name of the SQL Server

- 2000 clustered instance 114, 186, 438, 595, 848, 1060
- setting the internal name of the SQL Server
 - 2005 clustered instance 266, 341, 512, 713, 951, 1133
- switching online nodes (RDC) 651, 768
- verifying campus cluster configuration 449, 524
- verifying the HA failover configuration 125, 278
- verifying the primary site configuration for DR 815, 918
- configuration overview
 - SQL Server 2000
 - campus cluster 388
 - disaster recovery 785, 1009
 - high availability 57
 - replicated data cluster 545
 - SQL Server 2005
 - campus cluster 462
 - disaster recovery 888, 1082
 - high availability 207
 - replicated data cluster 662
- configure
 - LLT over UDP using VCW 92, 169, 243, 324, 405, 479, 565, 682, 805, 908
- converting standalone SQL Server to HA 133, 287
 - moving data files and databases 179, 334

D

- database
 - new database for SQL Server 2000 116, 440, 605, 1067
 - new database for SQL Server 2005 269, 515, 722, 1141
- disaster recovery (DR)
 - cloning secondary storage (VVR replication) 829, 931
 - cloning SQL Server 2000 service group 850
 - cloning SQL Server 2005 service group 954
 - configuration example 63, 212

- configuring DR components manually 1147
 - configuring GCO with DR wizard 854, 972
 - configuring replication with DR wizard 854, 958
 - creating temporary storage (array-based replication) 833, 935
 - defined 773
 - deploying for SQL Server 2000 777, 1001
 - deploying for SQL Server 2005 879, 1073
 - DR wizard overview 824, 927
 - DR wizard requirements 824, 927
 - illustrated 775
 - IP addresses 64, 215
 - multiple sites 874, 978
 - overview 773
 - setting up a secondary SQL Server 2000 site 1070
 - setting up a secondary SQL Server 2005 site 927, 1144
 - typical configuration 775
 - verifying HA configuration at primary site 815, 918
 - verifying the configuration 1069, 1143
 - disk groups
 - cloning for secondary site (array-based replication) 833, 935
 - cloning for secondary site (VVR replication) 829, 931
 - configuring for campus cluster 418, 492
 - creating
 - campus cluster 419, 493
 - DR 1042, 1115
 - HA 81, 157, 231, 312
 - MSDTC 359, 1176
 - RDC 576, 693
 - deporting
 - campus cluster 425, 431, 499, 505
 - DR 841, 945, 1053, 1127
 - HA 107, 180, 260, 335
 - RDC 588, 706
 - importing
 - campus cluster 425, 432, 499, 506
 - DR 842, 945, 1054, 1127
 - HA 108, 180, 260, 335
 - RDC 589, 707
 - multiple instances 79, 230
 - overview
 - campus cluster 416, 490
 - HA 77, 154, 228, 308
 - preconditions
 - campus cluster 416, 490
 - DR 1041, 1114
 - HA 78, 229
 - MSDTC 358, 1175
 - RDC 575, 692
 - sample configuration 78, 155, 229, 310, 417, 491
 - where to create 78, 154, 228, 309, 416, 490
 - DR wizard
 - cloning secondary storage
 - array-based replication 833, 935
 - VVR replication 829, 931
 - cloning SQL Server 2000 service group 850
 - cloning SQL Server 2005 service group 954
 - configuring replication and GCO 854, 958
 - overview 824, 927
 - requirements 824, 927
 - drive letters
 - adding drive letters to mount volumes
 - campus cluster 432, 506
 - adding to mount volumes
 - DR 842, 946, 1054, 1128
 - HA 108, 181, 261, 336
 - RDC 589, 707
 - driver signing options
 - resetting 227, 398, 472, 558, 675, 798, 901, 1024, 1097
- ## E
- EMC SRDF
 - configure SRDF replication with the DR wizard 862, 966
 - requirements for DR wizard 819, 922
- ## F
- failover considerations for SQL Server 2000 60
 - Fire Drill Wizard
 - actions 988
 - changing a fire drill configuration 995
 - deleting the configuration 998
 - overview 984
 - preparing the configuration 990
 - prerequisites for a fire drill 986
 - restoring the prepared configuration 996
 - running a fire drill 993
 - forceimport
 - attribute for campus cluster 374

defined 374
 setting after a site failure 449, 524

G

Global Cluster Option (GCO)
 configuring with the DR wizard 868, 972
 prerequisites 868, 972, 1162
 secure configuration 872, 976, 1167

H

hardware configuration for a cluster 68, 218, 547,
 664, 789, 892, 1012, 1085
 high availability (HA)
 adding HA to standalone SQL Servers 133, 287
 defined 45
 deploying for SQL Server 2000 47
 deploying for SQL Server 2005 197
 overview for SQL Server 45
 VCS SQL Server agent 46
 verifying the failover 125, 278
 Hitachi TrueCopy
 configure replication with the DR wizard 865,
 969
 requirements for DR wizard 821, 924

I

installing SFW HA
 campus cluster 392, 466
 DR 792, 895, 1015, 1088
 HA 71, 221
 RDC 550, 666
 installing SQL Server 2000
 before installing on a second node 106, 587,
 840, 1052
 first node 103, 584, 837, 1049
 multiple instances 102
 second node 110, 182, 434, 591, 844, 1056
 installing SQL Server 2005
 before installing on a second node 259, 706,
 944, 1126
 first node 254, 701, 939, 1122
 multiple instances 254
 second node 262, 337, 508, 709, 947, 1129
 installing SQL Server 2000
 before installing on a second node 430
 installing SQL Server 2005
 before installing on a second node 505

internal name
 setting for the SQL 2000 clustered
 instance 114, 186, 438, 595, 848, 1060
 setting for the SQL 2005 clustered
 instance 266, 341, 512, 713, 951, 1133

IP addresses
 active/active configuration 62, 212
 active/passive configuration 59, 209
 disaster recovery configuration 64, 215

L

LLT over UDP
 configuring using VCW 92, 169, 243, 324, 405,
 479, 565, 682, 805, 908

M

MSDTC
 configuration tasks for DR 1169
 configuration tasks for HA 353
 configuring the client 367, 1182
 creating an MSDTC service group 365, 1180
 DR configuration overview 1171
 HA configuration overview 355
 mounting drives 364, 1179
 service group configuration 355, 1171
 multiple DR sites 874, 978
 multiple SQL Server instances
 assigning the port 118, 189, 271
 disk groups 79, 230
 failover considerations for SQL Server 2000 60
 installation order 102, 254
 instance name 102, 254
 service groups 119, 272
 virtual server name 102, 254

N

network configuration for the cluster 68, 218, 389,
 463, 547, 664, 789, 892, 1012, 1085

P

port assignment for multiple instances 118, 189,
 271
 prerequisites
 MSDTC service group
 DR 1171
 HA 355
 service group for SQL Server 2000

- campus cluster 442
 - HA 120, 189
- service group for SQL Server 2005
 - campus cluster 517
 - HA 273, 344
- SQL Server 2000
 - DR 781, 1005
 - HA 53
 - RDC 541
 - standalone conversion 137
- SQL Server 2005
 - DR 883, 1077
 - HA 203
 - RDC 657
 - standalone conversion 291
- primary host, defined 773
- primary site
 - verifying the cluster configuration 815, 918
- primary system zone 602, 721

Q

Query Analyzer 114, 186, 438, 595, 848, 1060

R

- replicated data clusters
 - defined 529
 - deploying for SQL Server 2000 537
 - deploying for SQL Server 2005 653
 - overview of setup 532
 - primary system zone 602, 721
 - secondary system zone 608, 725
 - verifying the configuration 607, 724
- replication
 - configuring for EMC SRDF with the DR wizard 862, 966
 - configuring for HTC with the DR wizard 865, 969
 - configuring for VVR with DR wizard 854, 958
 - defined 774
 - setting up a Replicated Data Set (RDS) 617, 734, 1149
- resetting
 - driver signing options 227, 398, 472, 558, 675, 798, 901, 1024, 1097
- resources
 - adding to SQL 2000 service group
 - campus cluster 450
 - HA 131

- adding to SQL 2005 service group
 - campus cluster 525
 - HA 284

S

- sample configurations
 - SQL Server 2000
 - DR 787, 1011
 - HA active/active 61
 - HA active/passive 58
 - RDC 546
 - SQL Server 2005
 - DR 891, 1084
 - HA active/active 210
 - HA active/passive 209
 - RDC 662
 - standalone SQL Server 2000 conversion to
 - HA 142
 - standalone SQL Server 2005 conversion to
 - HA 297
- secondary host, defined 773
- secondary site
 - setting up for disaster recovery 927, 1144
- secondary zone
 - setting up for replicated data cluster 608, 725
- secure clusters
 - assigning user privileges 823, 926
- secure GCO, establishing 872, 976, 1167
- Security Services
 - configuring 93, 170, 244, 325, 406, 480, 566, 683, 806, 909, 1031, 1104
- service groups
 - configuring for SQL Server 2000 119, 597, 850, 1062
 - configuring for SQL Server 2005 272, 716, 954, 1136
 - dependencies 788, 875, 892, 979
 - modifying for SQL 2000 131, 450
 - modifying for SQL 2005 284, 525
 - prerequisites for configuring for SQL Server 2000 120, 189
 - prerequisites for configuring for SQL Server 2005 273, 344
 - priority order 119, 272
 - requirements for active/active configurations 119, 272
 - requirements for multiple instances 119, 272
- VCS SQL Configuration wizard 121, 191, 274, 345, 443, 518

- setting bandwidth
 - using RDS wizard 859, 963
- SFW HA installation
 - DR 792, 895, 1015, 1088
 - HA 71, 221
 - RDC 550, 666
- site failure, forceimport attribute 449, 524
- Solutions Configuration Center
 - context sensitivity 30
 - overview 29
 - running wizards remotely 37
 - starting 30
 - wizard descriptions 37
 - workflow for active/active configuration 67, 217
 - workflow for active/passive configuration 389, 463
- SQL Server 2000
 - before installing on a second node 106, 430, 587, 840, 1052
 - configuring a service group 442
 - configuring network and storage 68, 547, 789, 1012
 - configuring service group 119, 597, 850, 1062
 - creating a new database 116, 440, 605, 1067
 - DR configuration overview 545, 785, 1009
 - DR sample configuration 546, 787, 1011
 - HA configuration overview 57
 - HA sample configuration 58
 - installing first node
 - campus cluster 427
 - DR 837, 1049
 - HA 103
 - RDC 584
 - installing second node
 - campus cluster 434
 - DR 844, 1056
 - HA 110, 182, 434
 - RDC 591
 - IP resource, modifying 448
 - new database 116, 440
 - new installation tasks (DR) 777, 1001
 - new installation tasks (HA) 47
 - overview of user-defined database 605, 1067
 - removing shared files 110, 113, 182, 434, 437, 591, 844, 1056
 - setting services to manual start 106, 430, 587, 840, 1052
 - stopping the SQL Server service 107, 431, 588, 841, 1053
 - tasks for setting up a secondary site 1070
 - tasks for setting up an RDC secondary zone 608
 - user-defined database 116, 440
- SQL Server 2005
 - before installing on a second node 259, 505, 706, 944, 1126
 - configuring network and storage 218, 664, 892, 1085
 - configuring service group 272, 517, 716, 954, 1136
 - creating a new database 269, 515, 722, 1141
 - DR configuration overview 662, 888, 1082
 - DR sample configuration 662, 891, 1084
 - HA configuration overview 207
 - HA sample configuration 61, 209, 211
 - installing first node
 - campus cluster 501
 - DR 939, 1122
 - HA 254
 - RDC 701
 - installing second node
 - campus cluster 508
 - DR 947, 1129
 - HA 262, 337
 - RDC 709
 - Management Studio 266, 341, 512, 713, 951, 1133
 - new database 269, 515
 - new installation tasks (DR) 879, 1073
 - new installation tasks (HA) 197
 - overview of user-defined database 722, 1141
 - removing shared files 262, 337, 508, 708, 947, 1129
 - removing shared files on the second node 266, 341, 512
 - service group
 - IP resource, modifying 523
 - setting services to manual start 258, 504, 705, 943, 1125
 - start a SQL Server service 267, 341, 512, 713, 951, 1133
 - stopping the SQL Server service 259, 505, 706, 944, 1126
 - tasks for setting up a secondary site 1144
 - tasks for setting up an RDC secondary zone 725
 - user-defined database 269, 515

SQL Server Configuration Wizard 121, 191, 274, 345, 443, 518

SRDF

configuring replication with the DR wizard 862, 966
requirements for DR wizard 819, 922

standalone SQL Server

configuring network and storage 143, 298
converting to clustered for HA 133, 287
installing and configuring SFW HA 147
overview of HA configuration 141, 296
preparing for conversion (HA) 146, 300
sample HA configuration 142, 297

storage cloning with the DR wizard

for array-based replication 833, 935
for VVR replication 931

storage hardware configuration 68, 218, 547, 664, 789, 892, 1012, 1085

switching online nodes 651, 768

system zone 602, 721

T

tempdb database (campus cluster) 441, 516

tempdb database (HA) 117, 188, 270

U

user privilege assignment 823, 926

V

VCS

configuring the cluster 86, 237, 558, 676, 799, 902, 1024, 1097
configuring the SQL Server 2000 service group 119, 597, 850, 1062
configuring the SQL Server 2005 service group 272, 716, 1136
switching online nodes 651, 768

VCS Configuration Wizard 86, 163, 237, 318

VCS SQL Configuration Wizard 121, 191, 274, 345, 443, 518

virtual server name 102, 254

volumes

adding drive letters 108, 181, 261, 336, 432, 506, 589, 707, 842, 946, 1054, 1128
configuring for campus cluster 418, 492
considerations for VVR and disaster recovery 80, 156, 230, 311

creating on a cluster disk group 82, 159, 233, 314, 421, 495, 578, 695, 1044, 1117

dismounting 425, 499

mounting 425, 499

multiple instances 230

overview (campus cluster) 416, 490

preconditions on a cluster disk group 78, 229, 416, 490, 575, 692, 1041, 1114

preconditions on a cluster disk group (MSDTC) 358, 1175

sample configuration 78, 155, 229, 310, 417, 491

VVR

configuration diagram 63, 213, 786, 890, 1010, 1083

configuration diagram (MSDTC) 1174

configuration tasks 1147

configuring replication with DR wizard 854, 958

moving tempdb to separate volume 117, 188, 270, 441, 516

setting up RDS 617, 734, 1149

VxSAS 816, 919

VxSAS 816, 919

Z

zone 602, 721